Snapshot EFB

User Guide

Version 1.5

Contents

Windows Client Agent Installation 3

Caddy Handling 6

Immediate Manual File Backup 8

Immediate Image Backup 10

Creating Machine Groups 12

Creating an Automated Backup Schedule 15

File restore to point in time 17

Machine "Bare Metal" image restore 18

OS image restore to different physical hardware 19

Bare Metal Recovery of a Physical Windows Machine using PXE boot 22

Production Machine File Restore to VMWare Virtual Machine 25

Exchange Server Backup 30

SQL Server Backup 45

Backup AD Database Backup 48

AD Database Restore 49

PUSH Synchronisation of Appliance to Appliance 50

PULL Synchronisation of Appliance to Appliance 53

Installation of Snapshot EFB eMail Archiving 56

Windows Client Agent Installation

A Snapshot client installation will install a small service and two directory paths onto the Windows machine (\usr and \etc). When an agent is installed from the Snapshot EFB appliance it will be automatically matched with that appliance.

1. Login to the server with local administration rights
2. Open a Windows Explorer window
3. In the address bar connect to the Snapshot EFB appliance share **b_snapshot**
   If your snapshot server is 192.168.1.55 then type
   \\192.168.1.55\b_snapshot

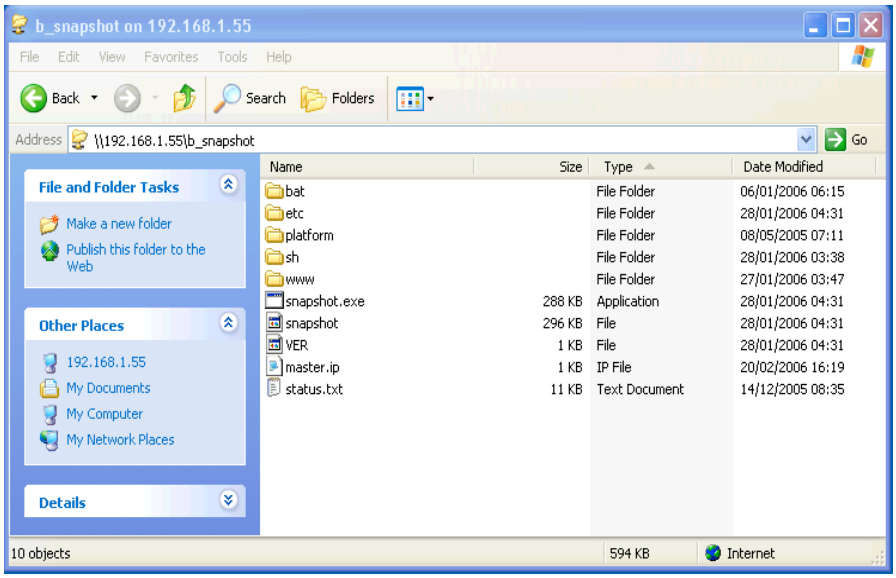

4. On the **b_snapshot** share navigate down the platform directory structure into the platform that you require. Note that Win is the 32bit windows agent and Win64 the 64bit Windows agent
5. Run the program **snapshot.exe** by double clicking on it. It will open up a command window
6. The agent installation program will ask if you ant to install snapshot type in **y**. The agent will install the Snapshot agent files and start the Snapshot EFB service on the Windows machine
7. The Snapshot agent is now installed on the Windows machine and ready to communicate to the Snapshot appliance it was install from

Now we need to force an agent update from the Snapshot appliance to compete the agent is install. This

process is usually done as a one off task against the whole of the machine estate and needs to be done from the Snapshot EFB appliance. You can either log into the Snapshot EFB appliance from a Windows machine using PUTTY or go physically from the Snapshot appliance itself. This procedure will use PUTTY from a Windows machine.

8. Start the PUTTY application and log in (see procedure Remote System Console)
9. Start the console by typing **setup_efb_server.sh** and pressing return. You will see the following menu

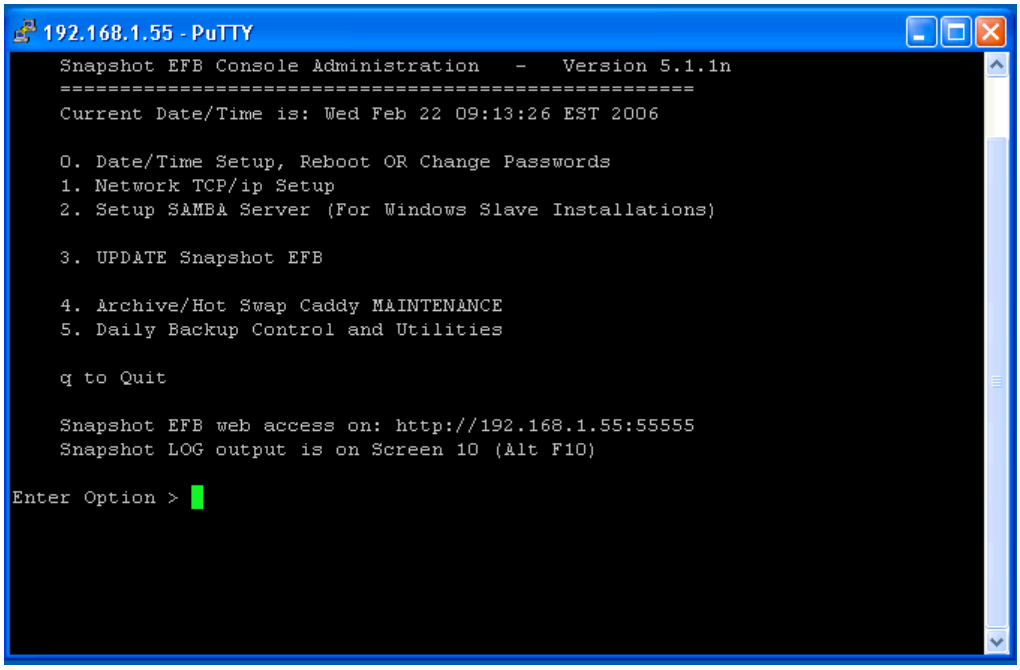
192.168.1.55 - PuTTY
Snapshot EFB Console Administration - Version 5.1.1n
Current Date/Time is: Wed Feb 22 09:13:26 EST 2006
0. Date/Time Setup, Reboot OR Change Passwords
1. Network TCP/ip Setup
2. Setup SAMBA Server (For Windows Slave Installations)
3. UPDATE Snapshot EFB
4. Archive/Hot Swap Caddy MAINTENANCE
5. Daily Backup Control and Utilities
q to Quit
Snapshot EFB web access on: http://192.168.1.55:55555
Snapshot LOG output is on Screen 10 (Alt F10)
Enter Option >

10. Select option **3 UPDATE Snapshot EFB**
11. Select option **3 AGENT Update / Poll from MASTER**
12. Enter in the IP address of the machine to be updated. Groups of machines can also be updated by typing in the group name (see procedure Creating Machine Groups)
13. The Snapshot EFB agent will be updated on all machines up and running and connected to the network
14. Hit return to return to the main menu
15. Select option **3 UPDATE Snapshot EFB**
16. Select option **4 Registry / SQL Export**
17. Enter in the IP address of the machine. Groups of machines can also be updated by typing in the

group name (see procedure Creating Machine Groups)

18. Snapshot will perform a registry export on all Windows servers.
19. Logout and closedown the PUTTY application

NOTE

For the agent installation to work there must be a local administrator account called **administrator** on the Windows machine the agent is being installed on.

Caddy Handling

As you are facing the Snapshot EFB appliance the caddies are the drives on the left hand side of the appliance with the silver handles. The pictuer below show two Snapshot EFB rack appliances.

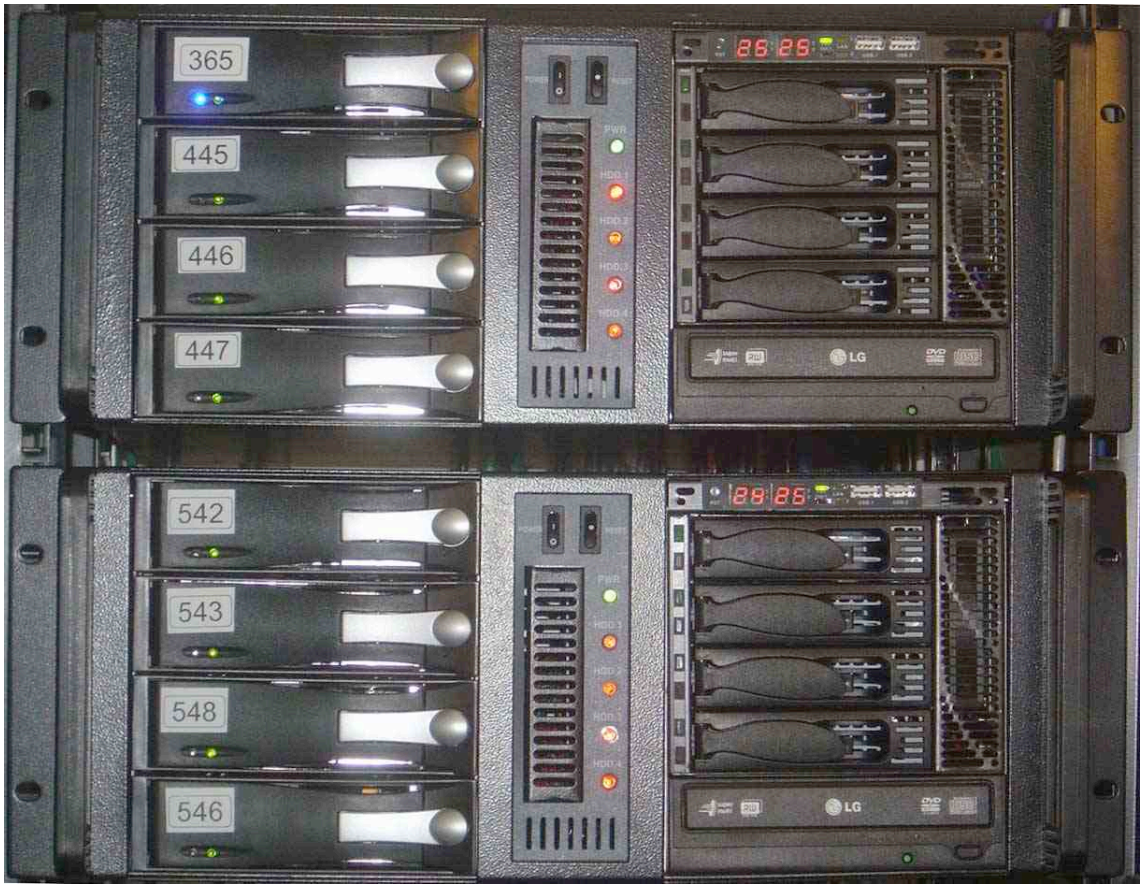

To change a caddy or set of caddies follow these instructions:

1. Snapshot EFB will have sent an email saying which caddies need to be replaced
2. The next set of caddies in the set will need to be requested from Iron Mountain to be delivered to the data centre. For example:
   Caddy 100-1a is replaced with 100-2a
   Caddy 100-2a is replaced with 100-3a
   Caddy 100-3a is replaced with 100-1a
3. Once the replacement caddy(s) have arrived, using the caddy key you will need to turn off ALL the caddies in the appliance. The key lock in under the silver handle
4. Remove the caddy that needs replacing and replace with the new caddy.

NOTE: If replacing multiple caddies then you will need to replace one caddy at a time, following steps 5 and 6 for each caddy.

***DO NOT* try and format a caddy with multiple caddies turned on.**

5. Turn on the new caddy using the caddy key
6. Format the new caddy
    a. Goto the Snapshot EFB console menu (white text based menu)
    b. Select menu option 4 caddy maintenance
    c. Select the option to format a new caddy
    d. The caddy may have and old data archive on it Snapshot will ask you if it is OK to remove it. Answer **yes** to this question
    e. The Caddy will be formatted. At the end of the formatting Snapshot EFB will ask you to give the caddy a name. It will suggest it's previous name that is the name on the label on the front of the caddy. Type in this name (e.g. 100-1a)
7. If you have to replace multiple caddies then turn off the caddy and go back to step 4 for the next caddy to be replace. If all caddies have been replaced and formatted then move onto step 8
8. Switch on all caddies
9. The next step is to assign an archive to the new caddy(s)
    a. Go to the Snapshot EFB console menu
    b. Select option 4 for caddy maintenance
    c. Select option 6 for assigning an archive to a caddy
    d. Snapshot EFB will go through the archives without a caddy assigned to them and ask you to enter which caddy you wish to assign to the archive.

    See section 2.5 for details on which archive goes to which caddy
10. The process is complete

Immediate Manual File Backup

An immediate manual file backup can be performed from the Snapshot browser interface. Either the administrator or the user of a machine, if an account for them has been setup, can instigate a manual file backup at anytime. However, it is normal for machine backups to be run in the background and automated.

A manual file backup can be performed whilst the host operating system, e.g. Windows, is up and running. There is no need for the machine to be logged out. Snapshot will automatically deal with locked files whether the files are locked by the operating system or user applications.

Where a file is altered whilst Snapshot is backing it up, Snapshot will detect this and skip the file. A warning will be logged against the backup indicating what happened and the file in question.

The following procedure will perform a complete file level backup of all non-removable hard discs.

1. Ensure that the machine is up and running Windows. It is better if the machine is not being used but this is not a requirement.

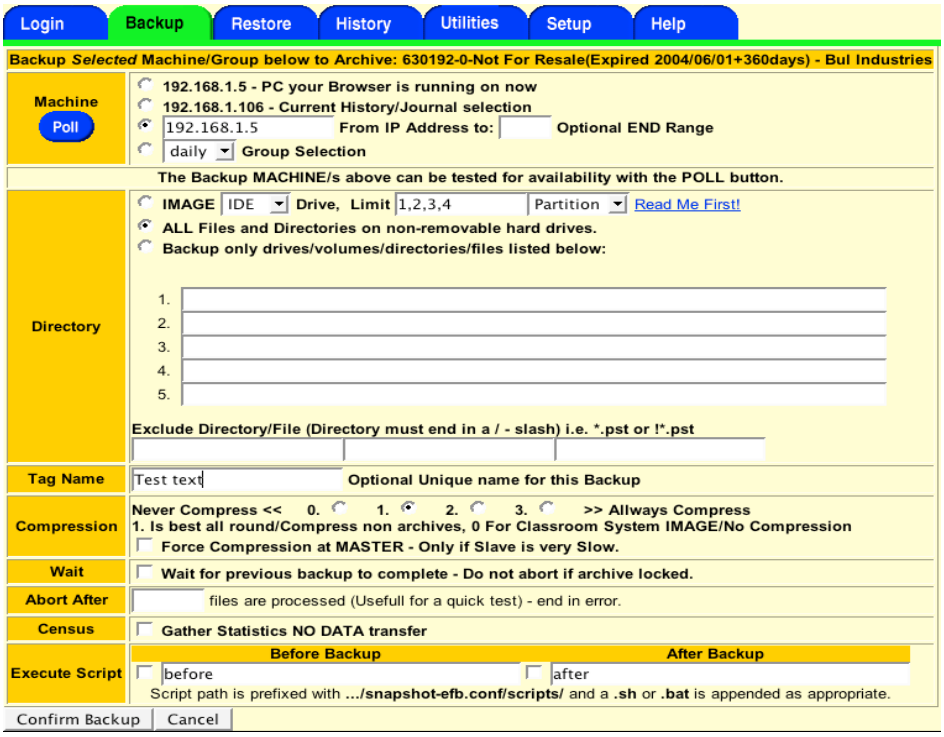

2. On the Snapshot console PC using the browser [Snapshot server IP:55555] :
    a. Login to the Snapshot archive which you wish to contain the backup
    b. Select the backup tab

    c. Type in the IP address of server waiting to be backed up or it's machine name
    d. Select the radio button "ALL Files and Directories on non-removable hard drives"
    e. In the "Tag Name" field type in a description of the backup
    f. Press the Confirm Backup button
    g. Press the Confirm button to start the backup

The browser will wait for the backup to complete and display a backup report page. The time this will take will depend on the number and size of files being backed up.

Immediate Image Backup

An immediate manual image backup can be performed from the Snapshot browser interface. Either the administrator or the user of a machine, if an account for them has been setup, can instigate a manual image backup at anytime. However, it is normal for machine backups to be run in the background and automated.

Just like with a file backup a manual image backup can be performed whilst the host operating system, e.g. Windows, is up and running. It is preferable for the machine to be logged out whilst the backup is being performed.

Altering files whilst Snapshot is backing up the machine may compromise the image backup. To ensure that this does not become a problem it is recommended, if possible, that an image backup is take whilst the machine operating system is not running.

The follow procedure takes an image backup of a machine booted from the Snapshot recovery CD and not running the host operating system.

1. The follow preparatory steps are recommended for the creation of a well compacted image but NOT mandatory:
    a. Copy cleanup.exe from CD or snapshot server share to c:\usr\snapshot-efb (if it does not already exist)
    b. Run cleanup.exe to clean unused sectors on hard disc to enable Snapshot to create a smaller image. This will take some time so be patient
        i. Start and command prompt and type:
        ii. cd \usr\snapshot-efb
        iii. cleanup c:
        iv. exit
    c. On CD or Snapshot server share run the ClearPageFileAtShutdown program. Do this from Windows explorer by locating the file is in the BAT directory and double clicking on it
    d. If possible defragment the machine
2. Reboot the server using the Snapshot CD or Snapshot PXE LAN boot if configured

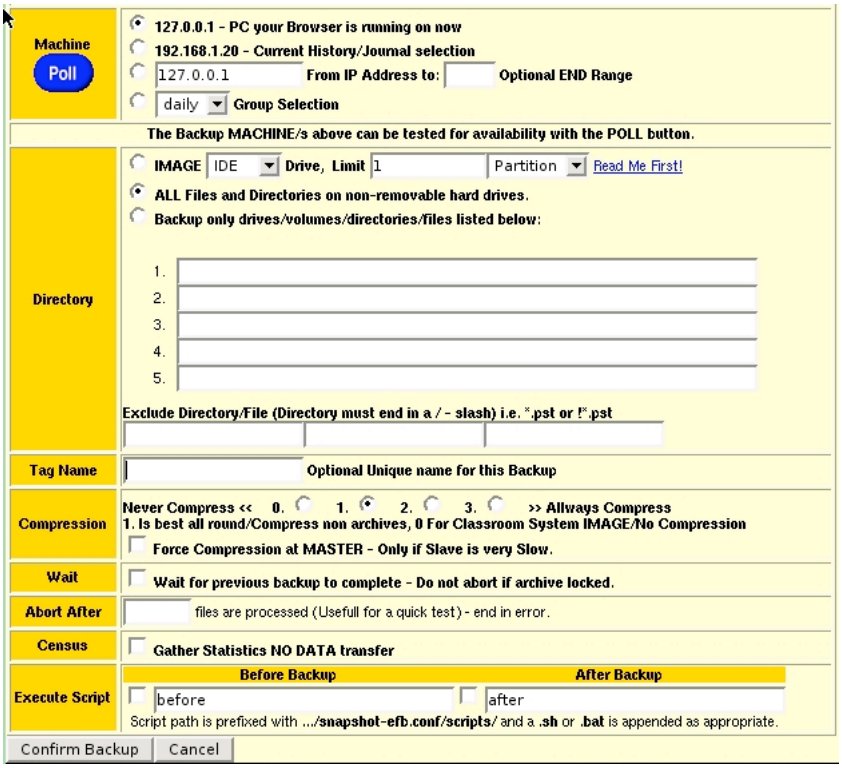
Machine
Poll
127.0.0.1 - PC your Browser is running on now
192.168.1.20 - Current History/Journal selection
127.0.0.1
From IP Address to:
Optional END Range
daily
Group Selection
The Backup MACHINE/s above can be tested for availability with the POLL button.
Directory
IMAGE
IDE
Drive, Limit
1
Partition
Read Me First!
ALL Files and Directories on non-removable hard drives.
Backup only drives/volumes/directories/files listed below:
1.
2.
3.
4.
5.
Exclude Directory/File (Directory must end in a / - slash) i.e. *.pst or !*.pst
Tag Name
Optional Unique name for this Backup
Compression
Never Compress << 0. 1. 2. 3. >> Allways Compress
1. Is best all round/Compress non archives, 0 For Classroom System IMAGE/No Compression
Force Compression at MASTER - Only if Slave is very Slow.
Wait
Wait for previous backup to complete - Do not abort if archive locked.
Abort After
files are processed (Usefull for a quick test) - end in error.
Census
Gather Statistics NO DATA transfer
Execute Script
Before Backup
After Backup
before
after
Script path is prefixed with .../snapshot-efb.conf/scripts/ and a .sh or .bat is appended as appropriate.
Confirm Backup
Cancel

3. On the Snapshot console PC using the browser (or any PC connected to the network):
    a. Login to the blank destination archive (if one does not exist then you will need to create one)
    b. Select backup tab
    c. Select IP address of server waiting to be imaged
    d. Select image
    e. SCSI or IDE depending on the machine being backed up
    f. Partition
    g. In Limit field type 1
    h. Enter in a tag name to name the backup
    i. Select the compression you want
    j. Confirm backup

The browser will wait for the backup to complete and display a backup report page. The time this will take will depend on the size of the partition being backed up, the speed of the network and the speed of the server being backed up (Typical 4Gb partition will take approximately 30 minutes to backup).

Please be patient and wait for the report page to be displayed.

Machines can also be booted from the network (PXE) or a USB key as well as CD if required.

Creating Machine Groups

This procedure will create machine groups to enable operations to be performed against multiple machines. For example, if you have some computers that are backed up only once a week a group can be created called "weekly" and a weekly backup operation scheduled against that group (see scheduled backups).

3. Write down the IP addressed or machine names of all the computers to be added to a group

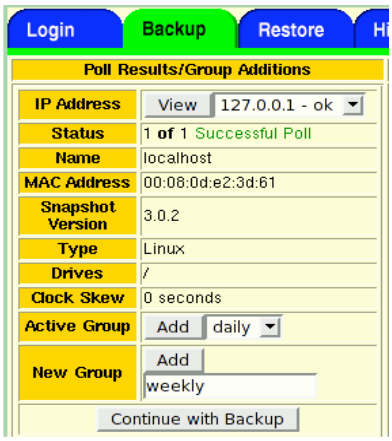

4. Group them together in logical groups e.g. Daily, Weekly, ArchiveA, ArchiveB etc
5. Login to the Snapshot Browser and select the Backup Tab. Click the POLL button and you will see the following Screen:

6. First use the "New Group" option to create your new group/s (by overtyping the weekly shown).
7. You may have to do this several times if your need multiple groups.
8. Now select the SETUP tab and from that screen select the "Slave" option. You will see the following Screen:

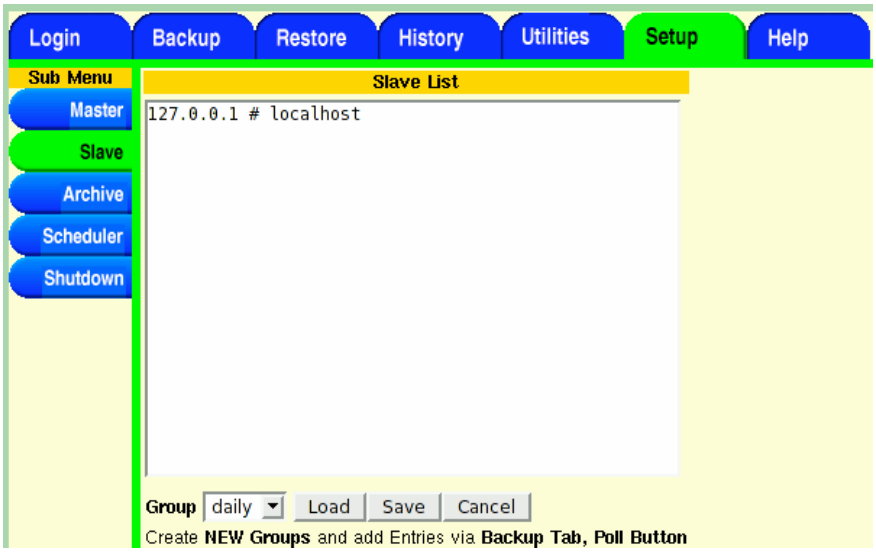

9. Select the Group you want to add individual computers to. The example above shows the "local host" to show you the "style" to be used.
10. Enter the IP addresses for this group. You may use the # to add computer descriptions.
11. Once you have added all the IP addressed for this group use the "Save" option to store the details away.
12. Do the same for other groups.
13. You can also use this option to remove or add new computers.
14. One completed press the "Scheduler" tab and you will see the following Screen:

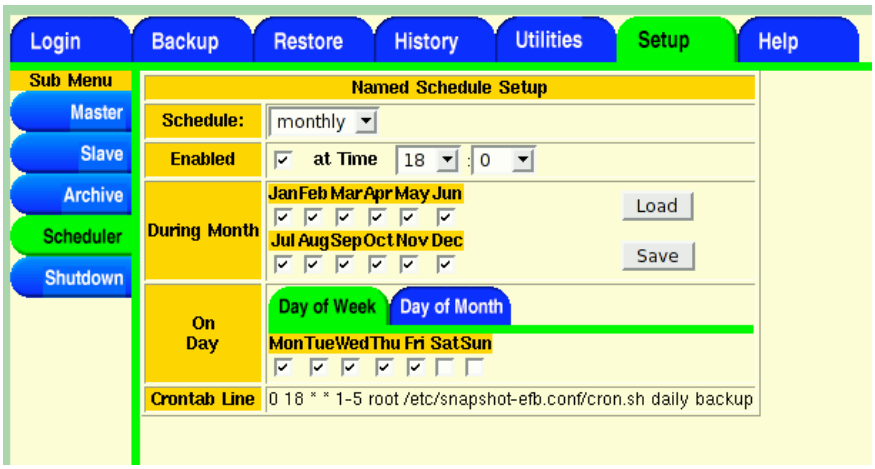

15. First select the group under the "Schedule" drop-down list.
16. Now Enable the option and enter the time (24 hour clock) at which time you want THIS group to be

backed up.

17. Now select the Day of the Month or Day of the Week for the backup to take place.
18. Use the "Save" button to create the entry in the "cron" file on the Snapshot Appliance so that the command as shown can be executed as setup by you.
19. Changes to the backup schedule can also be done on this screen, as a new "save" will over-ride the previous entry.

Creating an Automated Backup Schedule

Once an initial backup is done then it is normal practice to configure an automated scheduled backup. This is achieved through setting up a backup schedule (See Utility->Schedule). Backups are normally scheduled to run automatically at specific times on the Snapshot Appliance. Setting backups to run either daily, weekly or monthly is done using the Setup function in the Browser interface.

Before you can do this you must decide which computers are to be backed up together. Backup schedules are applied to machines groups. To setup a backup schedule you must first have setup a machine group to apply it to (see section Creating Machine Groups).

Backup schedules are applied to individual archives. Please note that if you have more than one archive configured on your appliance then you MUST be logged into the archive you are setting the schedule for. Schedules do not work across archives.

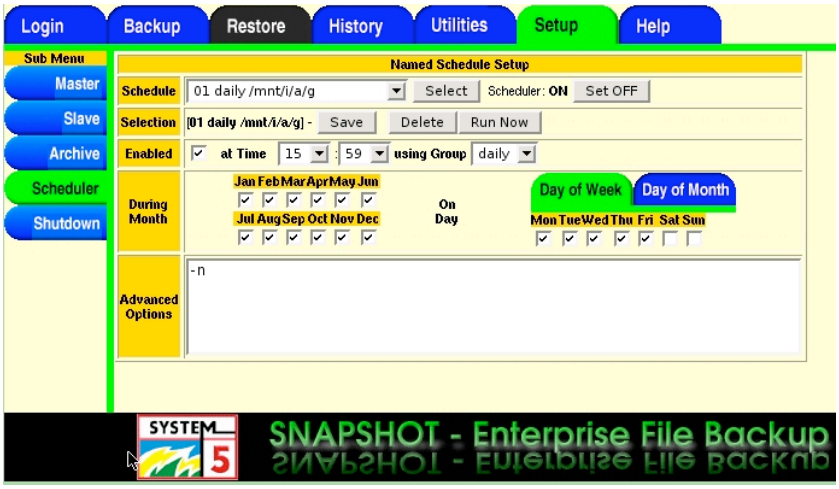

To create a schedule:

1. Using the Snapshot browser interface log into the archive you want to create a schedule for
2. Setup a machines group if you haven't already
3. Select the Setup tab across the top of the browser and the Scheduler tab on the left of the browser
4. Select the group you wish the schedule to apply to
5. Set the time the scheduled backup is to run at. The Snapshot Appliance uses a 24 hour clock
6. Set the months the schedule is to run in

7. Set the days of the week or the days in a month the schedule is to run on
8. Add any additional options you want to apply to the backup schedule to override the default backup operation (see appendix A for a list of common option commands)

   (NOTE: By default a backup schedule will perform a file backup of all files and directories on non-removable hard drives for all machines in the selected group)

   A good method of generating advanced options is to configure a manual backup using the browser interface. The interface will display the full backup command from which advanced options can be copied.

   For example to create a schedule to perform a Census add the command "-n" to the Advanced Options field (see diagram above).

9. Enable the schedule by checking the enabled checkbox
10. Save the schedule

By checking the initial archive usage after the first backup and the flow during the second and third you can estimate the capacity of the archive. Depending on the amount of data it usual to store 3 to 6 months of backups for all machines within an archive.

File restore to point in time

Use this procedure to recover a still bootable machine to a point in time using a Snapshot file backup. If you wish programs and setting to be fully recovered then a backup including a registry backup MUST be used. A full file backup will achieve this.

Any SQL server service should be stopped if their databases and configuration files are included in the recovery file set. If SQL server is not stop then certain key files may be locked preventing Snapshot from overwriting them.

1. Ensure that the machine is up and running Windows and sitting at the login prompt
2. On the Snapshot console PC using the browser enter the address [Snapshot server IP:55555] :
    a. Login to the Snapshot archive which contains the file backup you wish to restore
    b. Select the **History** tab
    c. Select the backup you wish to restore
    d. Select the **Restore** tab
    e. Select the appropriate **Restore To** radio button and type in the machine IP address or machine name of the machine you wish to restore to
    f. Select the **Original Directory (Path listed above under FROM:)** radio button
    g. Select the **Restore if any above OR MD5SUM fails** radio button
    a. Press the **Confirm Restore** button
    h. Press the **Confirm** button to start the restore

The browser will wait for the restore to complete and display a restore report page. The time this will take will depend on the number and size of files being restored. When the restore is complete log into the restored machine. If the backup contain a registry backup then the registry files will be restored and the machine rebooted.

NOTE:
Using this procedure Snapshot will not delete files that exist on the hard disc but do not exist in the backup being recovered. However, the registry, other configuration data and files that do exist within the backup will be restored back to their state when the backup was performed.

Machine “Bare Metal” image restore

Use this procedure to recover a machine from “bare metal” using an image backup. The advantage of an image restore over a file restore is an image restore can partition and format the hard disc as well as recover the files. A file restore expects that the drive is already partitioned and formatted.

To restore an image backup to a machine:

1. Reboot the server using the Snapshot CD or Snapshot PXE LAN boot if configured
2. On a machine using a browser enter the address [Snapshot server IP:55555] :
    a. Login to the Snapshot archive which contain the image you wish to restore
    b. Select history tab
    c. Select image
    d. Select restore tab
    e. Type in the IP address of target server waiting to be imaged
    f. Select image restore to same hard drive
    g. Confirm backup

The browser will wait for the restore to complete and display a restore report page. The time this will take will depend on the size of the partition being restored, the speed of the network and the speed of the server being restored (Typical 5Gb partition will take approximately 10 minutes to restore).

Please be patient and wait for the report page to be displayed.

Use this procedure in conjunction with a file based restore for data recovery and recovery a subsequent OS and application changes.

OS image restore to different physical hardware

Use this procedure to restore a machine image backup onto a different piece of hardware using different devices and components. This procedure will guide you through restoring the machine backup and then fixing the installation to work on the new hardware set.

NOTE:

For restoration of machines into Virtual Machines see the procedure backup verification.

This procedure will restore an image backup of a machine to a different hardware set. Use this procedure to recover a machine onto new hardware from "bare metal" creating a bootable server. A file backup instead of an image backup can also be used. If you choose to use a file backup then you will need to manually partition the drive(s) of the target machine.

For this procedure to work you will need a Windows OS installation disc and potentially a disc containing drivers for the new hardware set. The Windows OS installation disc MUST be the same version and revision level as the machine backup being restored.

To recover the machine backup:

1. Boot the new machine using the Snapshot CD or Snapshot PXE LAN boot if configured
2. On the Snapshot console or PC using the browser enter the address [Snapshot server IP:55555] :
    a. Login to the Snapshot archive which contain the image you wish to restore
    b. Select history tab

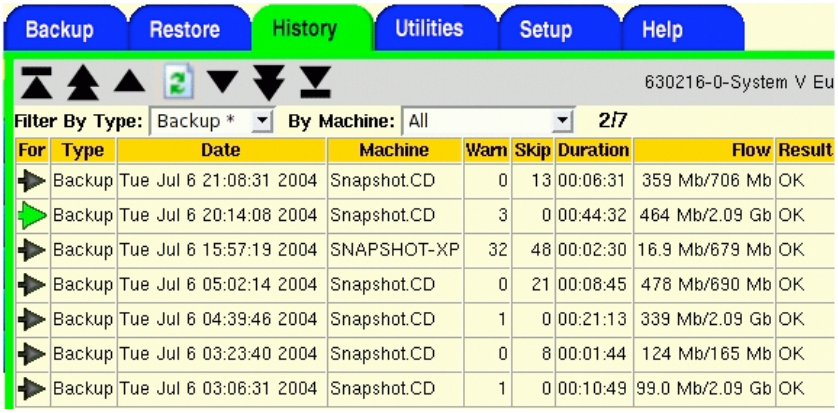

For	Type	Date	Machine	Warn	Skip	Duration	Flow	Result
	Backup	Tue Jul 6 21:08:31 2004	Snapshot.CD	0	13	00:06:31	359 Mb/706 Mb	OK
	Backup	Tue Jul 6 20:14:08 2004	Snapshot.CD	3	0	00:44:32	464 Mb/2.09 Gb	OK
	Backup	Tue Jul 6 15:57:19 2004	SNAPSHOT-XP	32	48	00:02:30	16.9 Mb/679 Mb	OK
	Backup	Tue Jul 6 05:02:14 2004	Snapshot.CD	0	21	00:08:45	478 Mb/690 Mb	OK
	Backup	Tue Jul 6 04:39:46 2004	Snapshot.CD	1	0	00:21:13	339 Mb/2.09 Gb	OK
	Backup	Tue Jul 6 03:23:40 2004	Snapshot.CD	0	8	00:01:44	124 Mb/165 Mb	OK
	Backup	Tue Jul 6 03:06:31 2004	Snapshot.CD	1	0	00:10:49	99.0 Mb/2.09 Gb	OK

c. Select image backup you wish to recover

d. Select restore tab

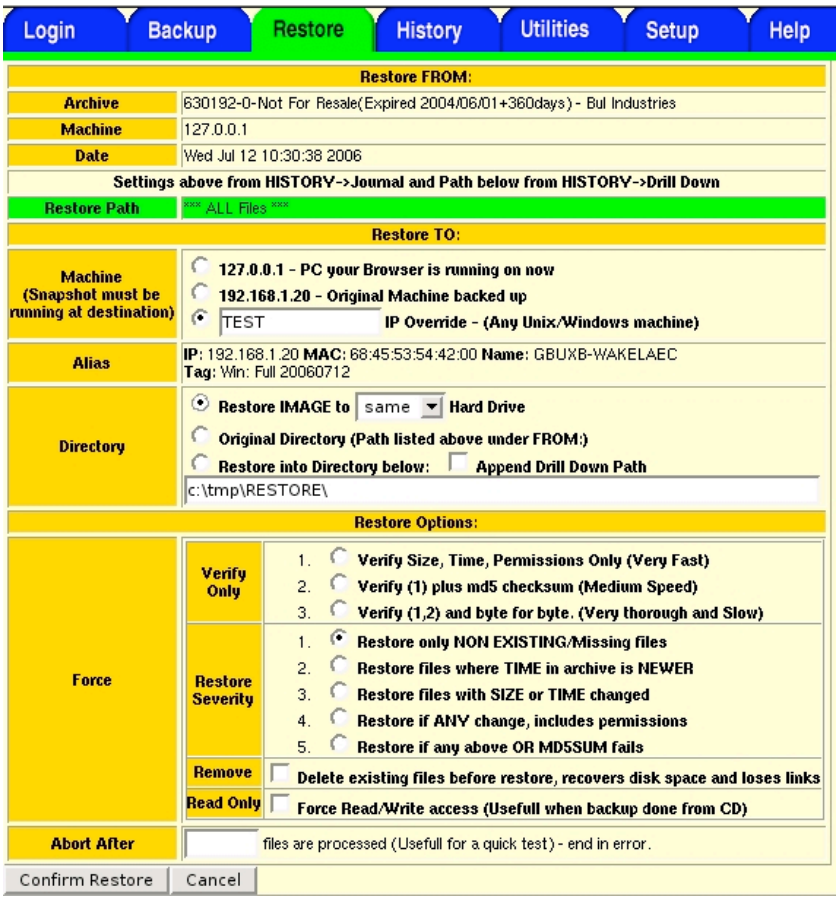
Login
Backup
Restore
History
Utilities
Setup
Help
Restore FROM:
Archive
630192-0-Not For Resale(Expired 2004/06/01+360days) - Bul Industries
Machine
127.0.0.1
Date
Wed Jul 12 10:30:38 2006
Settings above from HISTORY->Journal and Path below from HISTORY->Drill Down
Restore Path
*** ALL Files ***
Restore TO:
Machine (Snapshot must be running at destination)
127.0.0.1 - PC your Browser is running on now
192.168.1.20 - Original Machine backed up
TEST
IP Override - (Any Unix/Windows machine)
Alias
IP: 192.168.1.20 MAC: 68:45:53:54:42:00 Name: GBUXB-WAKELAEC
Tag: Win: Full 20060712
Directory
Restore IMAGE to same Hard Drive
Original Directory (Path listed above under FROM:)
Restore into Directory below:
Append Drill Down Path
c:\tmp\RESTORE\
Restore Options:
Force
Verify Only
1. Verify Size, Time, Permissions Only (Very Fast)
2. Verify (1) plus md5 checksum (Medium Speed)
3. Verify (1,2) and byte for byte. (Very thorough and Slow)
Restore Severity
1. Restore only NON EXISTING/Missing files
2. Restore files where TIME in archive is NEWER
3. Restore files with SIZE or TIME changed
4. Restore if ANY change, includes permissions
5. Restore if any above OR MD5SUM fails
Remove
Delete existing files before restore, recovers disk space and loses links
Read Only
Force Read/Write access (Usefull when backup done from CD)
Abort After
files are processed (Usefull for a quick test) - end in error.
Confirm Restore
Cancel

e. Type in the IP address or DNS name of target machine waiting to be imaged

f. Select image restore to same hard drive

g. Confirm backup

The browser will wait for the restore to complete and display a restore report page. The time this will take will depend on the size of the partition being restored, the speed of the network and the speed of the server being restored (Typical 5Gb partition will take approximately 10 minutes to restore).

Please be patient and wait for the report page to be displayed.

3. Reboot the machine. Hopefully, if the hardware is not too different, Windows will boot and fix itself up to the new hardware. If this happens then go to bullet 11 in this procedure
4. If the machine "Blue Screens" during boot then reboot machine using the Windows installation disc
5. The installation disc will load drivers, if it does not contain a SCSI driver for the new machine then load a new driver by pressing F6 when prompted
6. When prompted select install Windows. **Do not** go into the recovery console
7. The Windows installation will begin. It will discover an existing Windows partition with a Windows OS and ask if you want to repair it or install a new copy of Windows. Select the repair option. If this does not happen the chances are that you do not have a matching OS install disc to the machine backup you have recovered
8. Windows will now go through repairing the Snapshot recovered version of Windows. It will look like it is installing Windows, but all the programs and their settings will remain. New drivers will be installed for the new hardware
9. When recovery is finished boot into Windows
10. Perform any additional configuration changes required
11. Backup the new machine

Use this procedure in conjunction with a file based restore for data recovery and recovery a subsequent OS and application changes.

Bare Metal Recovery of a Physical Windows Machine using PXE boot

Use this procedure to recover a physical machine from bare metal using one of it's own full OS file level backups. Please note that the file backup that you use for this recovery must be a full OS backup containing all OS files.

Note: this procedure works only for Windows operating systems.

1) Start up the machine to allow for the OS drive (usually the C: drive) to be restored:

    a. Boot to the network using F12 (on some machines the boot options might be on a different function key), and record the MAC address.

2) Set the Snapshot server into PXE boot mode:

    a. Open the Snapshot Putty interface OR use the Snapshot main console
    b. Use option 6 on (PXE boot)
    c. Specify the MAC address recorded from the steps above
    d. Specify the IP address that the new VMware machine will be using
    e. Select option 2 for a 'win-pe' boot
    f. wait for the 2 services to start up
        i. DHCP
        ii. TFTP

        Make sure you look for any errors on service start up. If an error occurs then please refer to you Snapshot EFB support dealer.

3) Now that you have the MAC address you can reboot the physical machine and set to PXE boot form the network again:

    a. Restart the machine
    b. Select the boot manager by press F12 (the key may vary on some systems)
    c. Select to boot from the network

d. The machine will now connect via the MAC address to the Snapshot server and download a recovery OS and boot up

4) Prepare the machine for the restore process:

a. On the machine being recovered – If you need to repartition and / or format the machine (i.e. if it has a new hard drive) then do so now. Below is a set of instructions to run a script that will repartition the hard drive to be one C: partition, format it and set it to be the boot partition.

If you do not need to repartition or format the hard drive then go to point b

i. Open a command prompt, 'GO > Command Prompt'
ii. Enter 'cd \etc\snapshot-efb.conf\scripts'
iii. Enter 'Createc'
iv. Enter 'exit'

NOTE:
If you do not wish to have the whole drive partitioned and formatted as a C: drive then you can run the diskpart.exe and format commands yourself. Look at the Createc.bat script for an example of how to use these commands.

5) Restore the C drive:

a. Open up a Snapshot Mozilla interface
b. Restore only the C drive of the required machine from the selected backup to the IP address of the new machine and restoring to the original directory

6) Ensure that the C drive is the primary partition for booting from:

a. In the machine, open a new command prompt, 'GO > Command Prompt'
b. Go to the C drive, 'cd c:'
c. Enter 'attrib –s –h boot.ini' to make te file visible and editable
d. Open up 'GO > Programs > Accessories > Wordpad' and edit the C:\boot.ini file

e. If necessary update the partition number to (1) if set to (2) (two entries)
f. Save the updates and close wordpad
g. Enter 'attrib +s +h boot.ini' to restore the system file settings
h. Close the command prompt

Q & A

Question: What if I want to have a machine with a C: and D: partition?

Answer: Once you have booted the machine into the recovery OS you can partition and format the machines' hard drive however you like using the partitioning tool. For Windows this is diskpart.

Question: Why does Snapshot but the partitions back to how they were when the backup was taken?

Answer: If this case we are restoring a machine back at a file level. Snapshot will recover the files back to which ever drive and directory path we tell it to. For a file level bare metal restore you will need to tell Snapshot to put the files back to there original positions and make sure that you have created and formatted the correct partitions if required.

If you use Snapshot to recover an image then snapshot is able to put back the partitions as they were when the machine was imaged. This is because in this case Snapshot is working at a disc sector level and not at a file level. The penalties for using imaging are that it takes more storage and takes longer both to back and recover a machine.

Question: Which is the best recovery OS to use?

Answer: It depends on which OS you are recovering. For Windows servers people commonly use a the Windows recovery OS.

Question: The machine I am trying to PXE does not start to boot the recover OS I selected?

Answer: There are several reasons why this may happen:

- Check the Snapshot appliance console menu where you set up the PXE boot and look to see if Snapshot is reporting an error starting up the DHCP server. If it is then this can be down to:
    - Samba not having been configured on the Snapshot appliance. Using option 2 on the main console menu to do this

- An incorrect network subnet. Commonly people use an IP address like 192.168.0.1 with a subnet mask of 255.255.255.0. This is a class B IP address with a class C subnet mask. The DHCP server on the Snapshot appliance can not automatically cope with this mismatch of network classes and fails to start. A manual network subnet can be entered into the /root/sh/start_pxe.sh script to cope with this. Edit the start_pxe.sh script and replace the NETWORK= line with a manual subnet. So for the example given above NETWORK = 192.168.0.0

- Make sure that you have entered in the correct MAC address for the machine you are trying to PXE booting

Production Machine File Restore to VMWare Virtual Machine

Use this procedure to create a new virtual machine in VMWare Server and restore the C: drive of a physical machine into the virtual machine. The whole f the virtual drive is partitioned as a C: drive, a Snapshot file restore is used to recover all C: drive files, the production machine restore is then fixed up to run in the virtual machine.

Note: this procedure works only for Windows operating systems.

1) Create a new VMware machine specifying the new name, file location and operating system:

   a. Start VMware and select File > New > Virtual Machine
   b. Next
   c. Select Typical Configuration – Next
   d. Select the correct server operating system for the server you are going to be reloaded – Next
   e. Enter the new VMware machine name and the folder you wish to store the data files in (remember this can need several GB of space) – Next
   f. Select Use Bridged Networking – Next
   g. Select a disk capacity of at least 2GB larger than the amount of data to be restored into the C drive of the server you are creating, Do not select 'allocate now' and unless a very large amount do not select 'split into 2GB files' – Next

2) Start up the VMware machine to allow for the C drive to be restored:

a. If on VMware workstation, disable the Audio
b. Boot to the network using F12, and record the MAC address

3) Set the Snapshot server into PXE boot mode:

a. Open the Snapshot Putty interface
b. Use option 6 on (PXE boot)
c. Specify the MAC address recorded from the steps above
d. Specify the IP address that the new VMware machine will be using
e. Select option 2 for a 'win-pe' boot
f. wait for the 3 services to start up

4) Connect the new VMware system to the network:

a. Stop the VMware machine
b. Restart the VMware machine
c. When it displays the VMware logo select the screen and press F12 to boot from the network
d. The VMware machine will now connect via the MAC address to the Snapshot server and download an Linux OS and boot up

5) Prepare the VMware machine for the restore process:

a. In the VMware machine open a command prompt, 'GO > Command Prompt'
b. Enter 'cd \etc\snapshot-efb.conf\scripts'
c. Enter 'Createc'
d. Enter 'exit'

6) Restore the C drive image:

a. Open up a Snapshot Mozilla interface
b. Restore only the C drive of the required machine from the selected backup to the IP address of the new machine and restoring to the original directory

7) Ensure that the C drive is the primary partition for booting from:

a. In the VMware machine, open a new command prompt, 'GO > Command Prompt'
b. Go to the C drive, 'cd c:'
c. Enter 'attrib –s –h boot.ini' to make te file visible and editable
d. Open up 'GO > Programs > Accessories > Wordpad' and edit the C:\boot.ini file
e. If necessary update the partition number to (1) if set to (2) (two entries)
f. Save the updates and close wordpad
g. Enter 'attrib +s +h boot.ini' to restore the system file settings
h. Close the command prompt

8) Set the C drive as a bootable partition with the correct drivers for the VMware environment:

a. Select 'GO > Programs > FIX_VMSCSI > Set OS target Root (current=)
b. Select the C:\Windows folder or C:\WINNT folder depending on the operating system.
c. Select GO > Programs > FIX_VMSCSI > (P2V) add VMSCSI (Buslogic) Windows XP/2000/2003
d. Shutdown the VMware system

9) Prepare the VMware machine for use:

a. Add any extra disks, D or E drives as necessary
b. Disable the Ethernet bridged network from starting up at power on (otherwise it will start up on the network and conflict with the live server)
c. Start up the VMware machine and let it boot into Windows
d. Sign on using the Domain Administrator Userid
e. Install the VMware Tools from the VM toolbar menu option
f. Update the IP address on the network card with the new IP address from above and the other DNS, Wins and Default gateway as necessary
g. Add a new Local Administrator Userid
h. Update the server name to the new VMware server name and make the server part of a local WORKGROUP not in the domain
i. Shutdown the VMware machine

10) Start up of the VMware machine:

a. Restart the VMware machine and let it boot up into windows
b. Sign on to the system using the new Local Administrator Userid
c. Add the VMware system to the domain
d. Use Disk Management to configure the new disks and the drive letters according to how the original system was configured
e. Restore the other drives using the Snapshot Mozilla interface
f. Restart the VMware system.

The system is now up and running as an image of the original system, but with a new name and IP address.

Q & A

Question: My Windows server blue screens when I reboot it?

Answer: Make sure you have completed to P2V steps in the procedure above and that you have chosen the correct drivers for your OS and version of VMWare.

Question: My machine says that it cannot find a boot partition?

Answer: Reboot using PXE and check the c:\boot.ini file. Makes sure that the partition number is set correctly. Often physical machines have hidden partitions that move the boot partition from partition 1 to partition 2. In this case you will need to set the partition entries in boot.ini to partition 1.

Question: My recovered Windows machine boots to the login screen, but when I try to login it tells me that I need to activate Windows?

Answer: This is a result of the way Microsoft does its licensing. Unless you have a Volume License for Windows then you will need to activate it again because too much of the hardware has changed and it is therefore a MS requirement to activate. Depending on the state of your VM and whether it can get Internet connectivity, you may have to ring the activation telephone number to activate your machine.

Call Microsoft, tell them that you are doing disaster recovery and they will give you an activation code.

Exchange Server Backup

By their very nature Exchange systems are complex and can be configured in many ways. As such it is not the intention of this manual to give a definitive statement on how to backup or recover your Exchange server. However, our experience shows that the following procedure works well for backing up an Exchange server in readiness for either a data store recovery, or a full "bare metal" recovery of an Exchange server.

The process for backing up an Exchange server is split into two parts, the initial backup and the ongoing backups. It assumes that the OS & Applications are stored on a separate partition to the Exchange database and logs.

1. **The initial Exchange server backup**

   It is important that we take a complete backup of the OS & applications partition and of the Exchange directory structure so we are able to easily recover the Exchange from "bare metal". Ideally the Exchange server would be a clean install at this point to minimise the amount of data we back up from the Exchange directory structure. Essentially we are going to do an offline Exchange server backup.

   a. Ensure that the Exchange Services are stopped. If you fail to do this files maybe altered whilst Snapshot is backing them up. This will cause snapshot to report a warning and skip the affected file.
   b. Follow the "Immediate Manual File Backup" procedure as outlined earlier in this manual
   c. Restart the Exchange services

2. **Ongoing Exchange Server Backup**

   This process falls into two parts, both of which are performed whilst the Exchange server is up and running.

   a. Online Data Store Export

      Backing up of the Exchange data directly out of the Exchange database. Exchange comes with a version of NTBackup that is tightly integrate into Exchange and makes it an ideal tool for doing online backups. NTbackup can perform full and differential backups of the Exchange system. We recommend that a full backup is performed once a week and differentials for the remainder of the week. This will provide the best balance between backup and recovery speed.

Use NTBackup to schedule:

A weekly full export into an export directory

A daily differential into an export directory

Performing an Exchange Information Store Online Backup

To perform an Online backup of Exchange Server 2003 follow these steps (note that 99% of the steps outlined here also work for Exchange 2000, although the screenshots were taken on Exchange Server 2003):

1. Run NTBACKUP
2. Click Advanced Mode

Backup or Restore Wizard
Welcome to the Backup or Restore Wizard
This wizard helps you back up or restore the files and settings on your computer.
If you prefer, you can switch to Advanced Mode to change the settings used for backup or restore. This option is recommended for advanced users only.
Always start in wizard mode
To continue, click Next.
< Back
Next >
Cancel

3. Click on the Backup tab.
4. On the Backup tab expand Microsoft Exchange Server, expand server name expand Microsoft Information Store, select the correct Storage Group's check-box or the entire Information Store check-box.

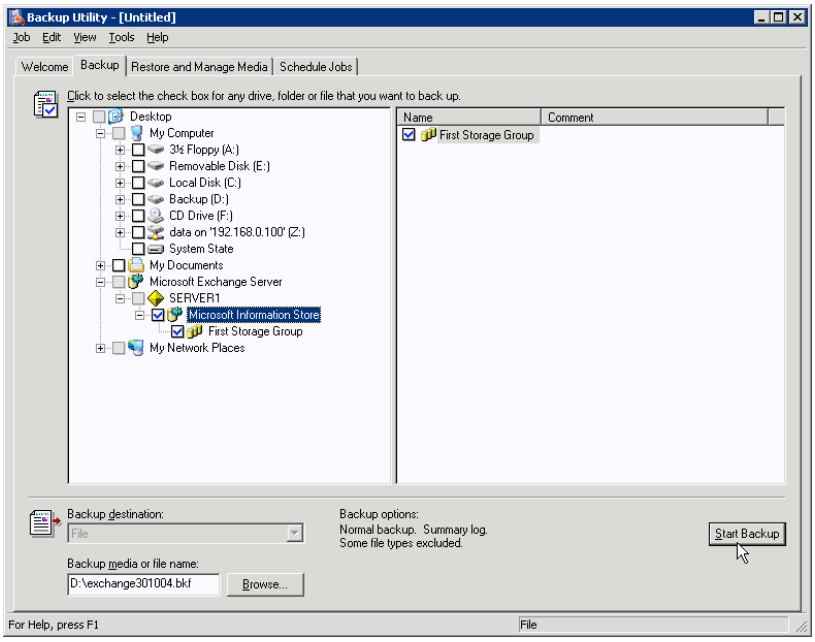

Note: It is always best to backup an entire Storage Group rather than individual Stores, and if possible, try to backup the entire IS in one job, although lack of time or storage might force you to selectively backup just parts of the IS at a time.

5. On the Backup choose to backup to a File.
6. In the Backup Filename make sure you enter a descriptive name for the backup job
7. Click on the Start Backup button.
8. In the Backup Job Information window make sure you have your settings right.

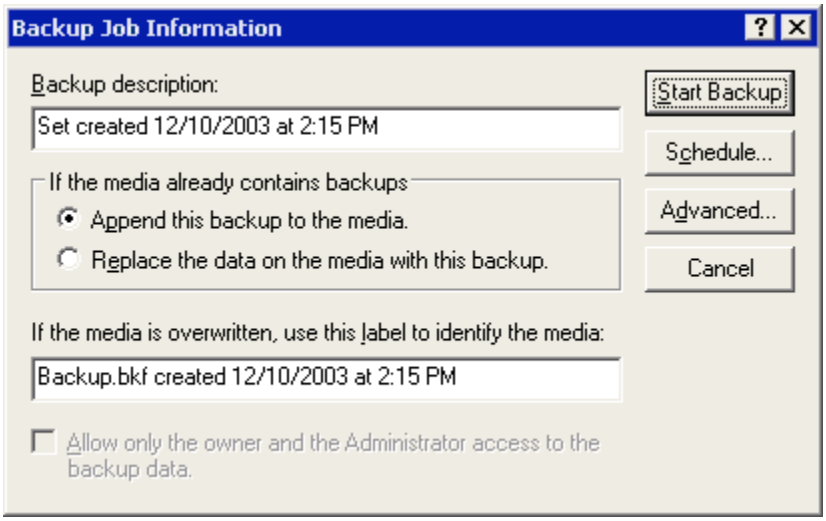
Backup Job Information
Backup description:
Set created 12/10/2003 at 2:15 PM
If the media already contains backups
Append this backup to the media.
Replace the data on the media with this backup.
If the media is overwritten, use this label to identify the media:
Backup.bkf created 12/10/2003 at 2:15 PM
Allow only the owner and the Administrator access to the backup data.
Start Backup
Schedule...
Advanced...
Cancel

Backup options	Explanation
Backup description	You can type a unique name to describe each backup.
Append this backup to the media.	You can append the backup job to an existing backup file or tape. Selecting this option does not affect any previous backup data stored on that media.
Replace the data on the media with this backup.	You can erase the backup file or all previous backup jobs stored on the destination media before the new backup job is saved. **Caution:** If you select this option, you will lose all previous backup data currently stored on the destination media.
Allow only the owner and the Administrator access to the backup data.	You can specify that only the owner or members of the Administrator's group may access the data that is saved on the destination media. This option secures the online tape or file. If you are backing up data to an existing tape or file that you want to overwrite, you can select this option. If you are backing up data to an existing tape or file and you are appending the data to the tape or file, you cannot select this option because ownership of the tape has already been established.

To schedule this backup to run automatically at a later time, in Backup Job Information, click Schedule.

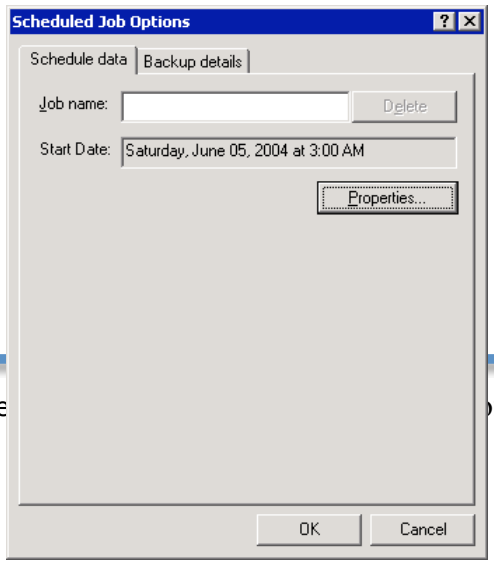

In Scheduled Job Options, in the Job name box, type a name for the scheduled backup job. Click Properties to enter the date, time, and frequency parameters for the scheduled backup, and then click OK

9. Now click on the Advanced button

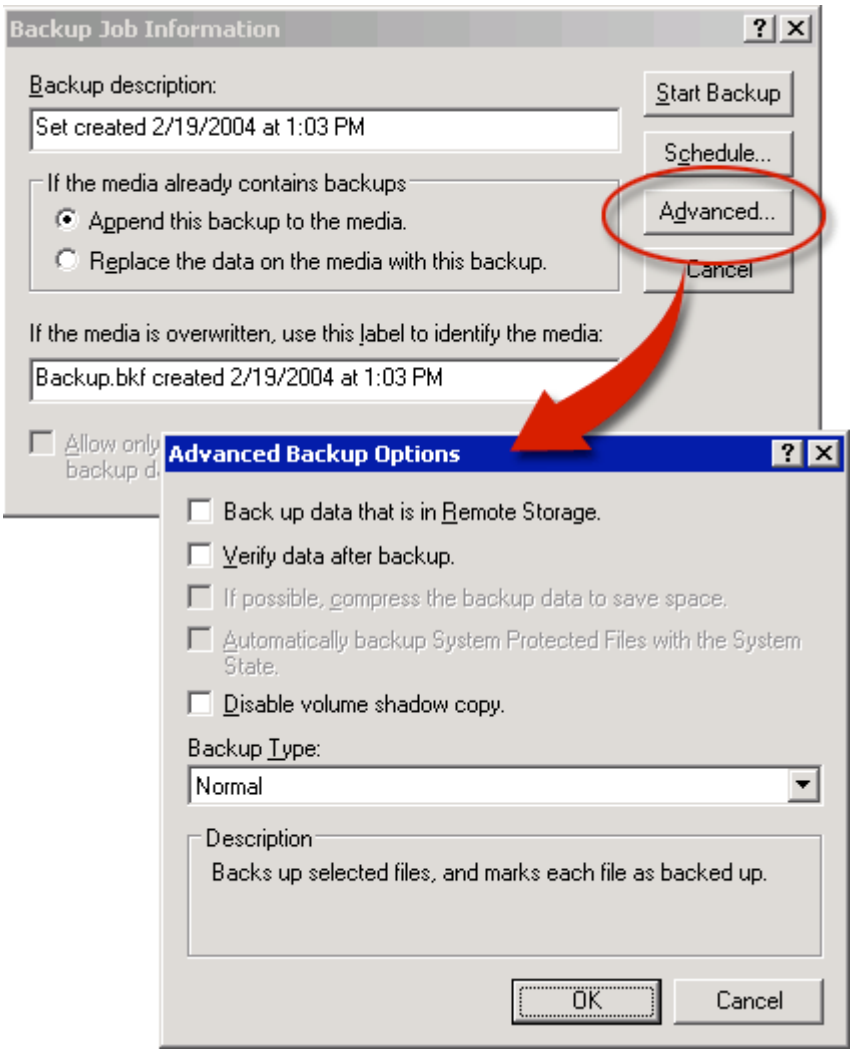
Backup Job Information
Backup description:
Set created 2/19/2004 at 1:03 PM
Start Backup
Schedule...
If the media already contains backups
Append this backup to the media.
Replace the data on the media with this backup.
Advanced...
Cancel
If the media is overwritten, use this label to identify the media:
Backup.bkf created 2/19/2004 at 1:03 PM
Advanced Backup Options
Back up data that is in Remote Storage.
Verify data after backup.
If possible, compress the backup data to save space.
Automatically backup System Protected Files with the System State.
Disable volume shadow copy.
Backup Type:
Normal
Description
Backs up selected files, and marks each file as backed up.
OK
Cancel

You can specify how your data is backed up. For example, you cannot perform an incremental backup of an Exchange store until you have performed a normal backup at least one time before the incremental backup.

Normal A normal backup entails copying all the files that you selected, and then marking each file as having been backed up. (The **archive** attribute is cleared.) To perform a normal backup, you need only the most recent copy of the backup file or

tape to restore all the files. You usually perform a normal backup the first time that you create a backup set. Normal backups are frequently referred to as full backups, although a backup is only a full backup when you choose to back up all files.

Important:
To restore from backup, you must have the most recent normal, differential, and incremental backup files.

Copy A copy backup entails copying all the files that you select, but not marking each file as having been backed up. (The **archive** attribute is not cleared.) To back up files between normal and incremental backups, perform a copy backup because it does not affect other backup operations.

Differential A differential backup entails copying files created or changed since the last normal or incremental backup, but not marking files as having been backed up. (The **archive** attribute is not cleared.) To perform a combination of normal and differential backups, you must have the last normal backup set and also the last differential backup set.

Incremental An incremental backup entails backing up only those files created or changed since the last normal or incremental backup, and then marking files as having been backed up. (The **archive** attribute is cleared.) If you use a combination of normal and incremental backups, you must have the last normal backup set and also all incremental backup sets to restore your data.

Daily A daily backup entails copying all the files that you selected and that are also modified on the same day that you perform the daily backup. The backed up files are not marked as having been backed up. (The **archive** attribute is not cleared.)
(It is common practice to perform a full backup on day 1 followed by 6 differential backups.)

Suggestion
Create two NTBackup schedules:

- *A full Exchange database export using the normal backup type that run weekly and is set to replace the data on the media when the backup is done*
- *A differential database backup using type differential that is set to run daily and overwrite the existing backup file*

10. The NTBackup part of the Exchange Online Information Store backup is completed.

11. Now on the Snapshot EFB appliance create a scheduled that backs up the directory the Exchange Information Store backup was written to. This will copy the backup to the Snapshot EFB appliance.

    Schedule the Snapshot EFB job leaving enough time for the Exchange Information Store export job to complete. The tine required will depend on the size of the Exchange Information Store and the type of backup performed.

Using NTBackup to backup Exchange 2007 on Windows 2008

Windows 2008 does not ship with a backup program that is capable of backing up Exchange 2007 data stores running on Windows 2008. However, it is possible to use NTBackup from Windows 2003 to perform a backup in the same way that NTBackup under Windows 2003 worked. This method even purges your transaction logs properly so both full and differential database exports can be performed.

The procedure for getting NTBackup working on a Windows 2008 Exchange 2007 server is as follows:

1. Load up a Windows 2003 x64 system. The x64 part is important, as you should be running Exchange on a x64 platform (unless you are using an Itanium), and you must use matching binaries for this. I used a Virtual Machine for this.
2. Copy the ntbackup.exe, ntmsapi.dll, and vssapi.dll from C:\windows\system32 into a new folder.
3. Burn the new folder you created containing the copied files to a CD, or copy it to USB drive. You can also copy it over the network to your Windows 2008/Exchange 2007 Server to skip the next step.

4. Copy the folder from the CD or USB drive to your Windows 2008/Exchange2007 server. Say C:\NTbackup\
5. Launch Ntbackup.exe. Select the Information Stores as you would normally under NTBackup and start your backup.

   After it completes you should have a quick and easy Exchange 2007 Backup on Windows 2008.

b. Individual Mailbox Export

An Exchange mailbox backup, often called a brick level mailbox backup is an approach where each mailbox on the Exchange server has its contents to a file on a disc, usually a PST file. Each mailbox is backed up individually into its own PST file making restoring an individual mailbox very easy.

Please keep in mind that a mailbox back approach is no replacement for full backup of the Exchange server and should only be used in conjunction with a full server backup. A mailbox backup will not be able to recover the Exchange information store in the event of a server failure.

Mailbox backup relies on MAPI to access each mailbox to re-create all of the mailbox data in the message store. Be aware that this can create a large amount of data to backup. However, a mailbox backup can be invaluable to quickly restore a users mailbox, the alternative being to recover the whole Exchange Information Store to a recovery Exchange instance. Not a quick or easy piece of work.

You can still do mailbox backups by using a tool from Microsoft called EXMERGE. Using EXMERGE you can backup any mailbox on any server and schedule it to run at a specific interval.

EXMERGE works with Exchange 5.5, 2000 and 2003. You must make sure that you obtain the correct copy of EXMERGE for the version of Exchange you are using.
Use Exmerge.exe to extract individual mailboxes into PST files on a daily bases. This will give you the ability to easily backup and recover individual mailboxes.

After downloading the correct version of EXMERGE, extract the files found in it by using tools like WinZip or WinRAR. You will now need to copy EXMERGE.EXE to the %Program Files%\Exchsrvr\bin folder of your Exchange server.

Next, you need to configure your user account to have full mailbox rights for the specific mailbox/mailboxes that you want to open. On Exchange 2000/2003 the Exchange Full Administrator permissions does NOT, by default, allow you to open any other user's mailbox.

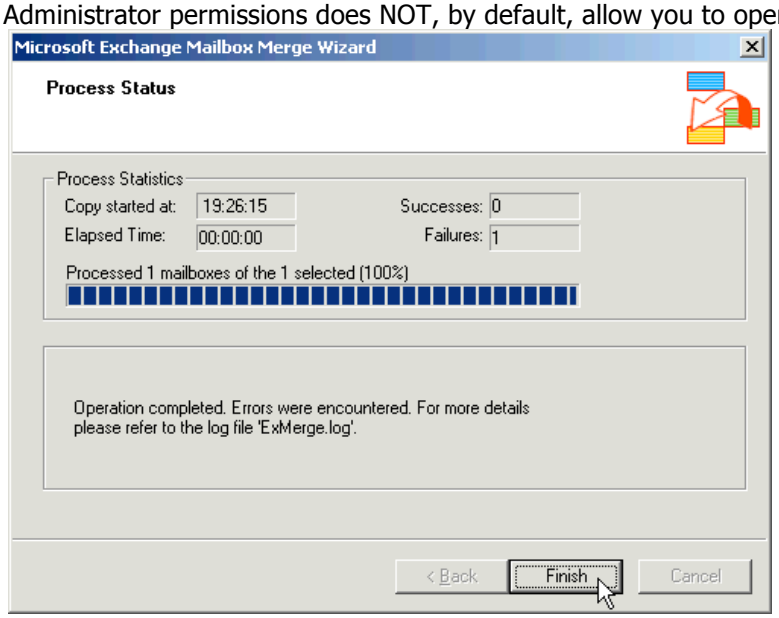

The above screenshot is an example of EXMERGE failure due to poorly configured permissions.

EXMERGE to Backup

To perform Brick-Level backups of one or more mailboxes found on one Exchange server follow these steps:

1. Navigate to the %Program Files%\Exchsrvr\bin folder of your Exchange server and double-click EXMERGE.EXE.
2. On the Welcome page click Next.

Microsoft Exchange Mailbox Merge Wizard
Welcome to the Microsoft Exchange Mailbox Merge Wizard
This program will allow you to extract data from mailboxes on a Microsoft Exchange Server and then merge this data into the same mailboxes on another Microsoft Exchange Server.
The program copies data from the source server into Personal Folders (.PST files) and then merges the data, in the Personal Folders, into the mailboxes on the destination server.
This version of the program will merge user folders and messages. It will merge server based folder rules only on servers running Exchange Server 5.0 or later. It does not support forms and Schedule+ data.
Messages copied will lose their single instance storage and this may result in the size of the Microsoft Exchange Private Information Storage increasing significantly.
< Back
Next >
Cancel

3. On the "Procedure Selection" dialogue box, select "Export or Import (Two Step Procedure)".

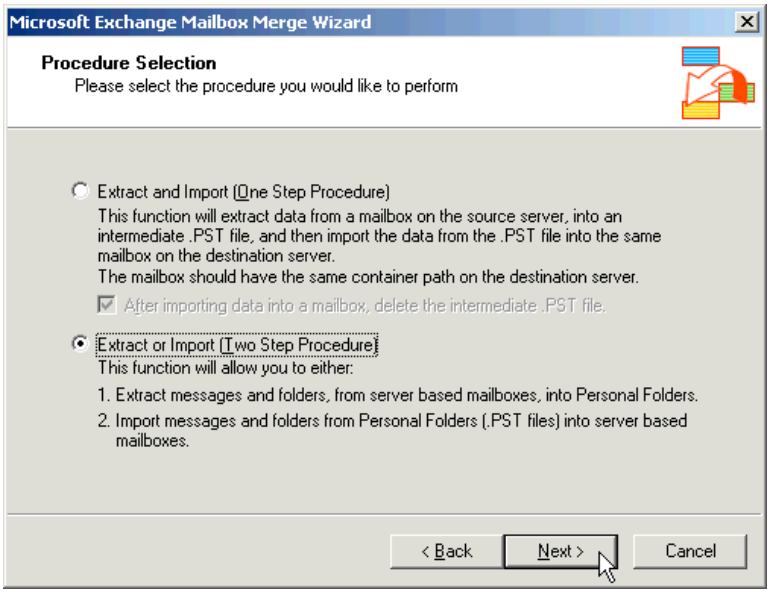
Microsoft Exchange Mailbox Merge Wizard
Procedure Selection
Please select the procedure you would like to perform
Extract and Import (One Step Procedure)
This function will extract data from a mailbox on the source server, into an intermediate .PST file, and then import the data from the .PST file into the same mailbox on the destination server.
The mailbox should have the same container path on the destination server.
After importing data into a mailbox, delete the intermediate .PST file.
Extract or Import (Two Step Procedure)
This function will allow you to either:
1. Extract messages and folders, from server based mailboxes, into Personal Folders.
2. Import messages and folders from Personal Folders (.PST files) into server based mailboxes.
< Back
Next >
Cancel

4. On the "Two Step Procedure" dialogue box, select "Step 1: Extract data from and Exchange Server Mailbox".

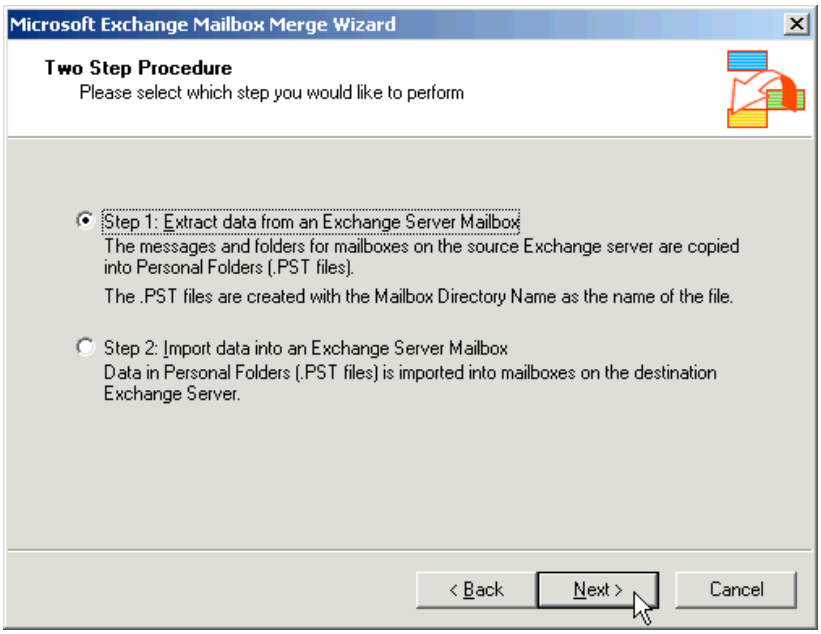
Microsoft Exchange Mailbox Merge Wizard
Two Step Procedure
Please select which step you would like to perform
Step 1: Extract data from an Exchange Server Mailbox
The messages and folders for mailboxes on the source Exchange server are copied into Personal Folders (.PST files).
The .PST files are created with the Mailbox Directory Name as the name of the file.
Step 2: Import data into an Exchange Server Mailbox
Data in Personal Folders (.PST files) is imported into mailboxes on the destination Exchange Server.
< Back
Next >
Cancel

5. In the "Source Server" dialogue box specify the name of your Exchange server. If you have a multiple domain environment you'll need to specify the name and LDAP port number of your Domain Controller.

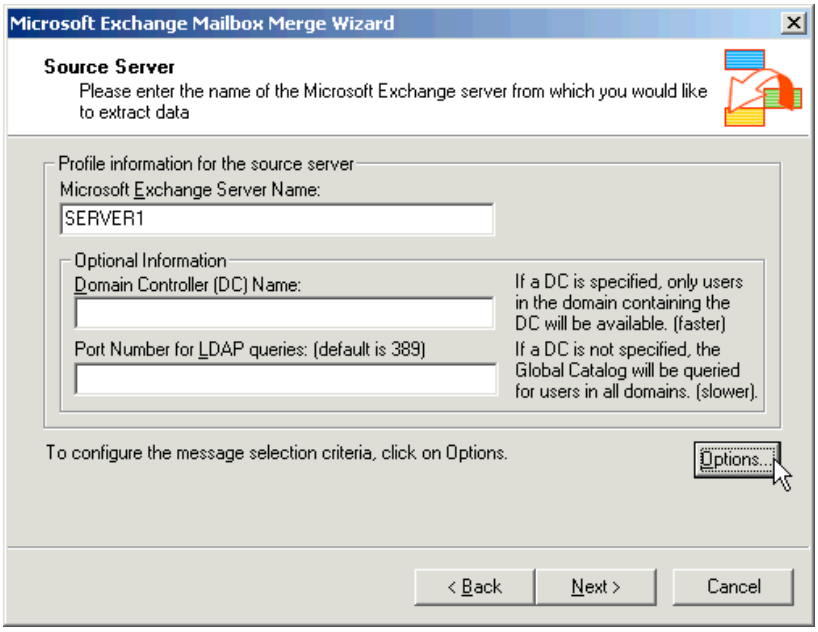
Microsoft Exchange Mailbox Merge Wizard
Source Server
Please enter the name of the Microsoft Exchange server from which you would like to extract data
Profile information for the source server
Microsoft Exchange Server Name:
SERVER1
Optional Information
Domain Controller (DC) Name:
If a DC is specified, only users in the domain containing the DC will be available. (faster)
Port Number for LDAP queries: (default is 389)
If a DC is not specified, the Global Catalog will be queried for users in all domains. (slower).
To configure the message selection criteria, click on Options.
Options...
< Back
Next >
Cancel

Click Next.

6. In the "Database Selection" dialogue box select the mailbox store you want to perform the action upon.

   Click Next.

7. In the "Mailbox Selection" dialogue box select individual mailboxes or press the "Select All" button to select all mailboxes found in that store. In this example I will only select one mailbox. Note that you can also see the mailbox size next to the mailbox name.

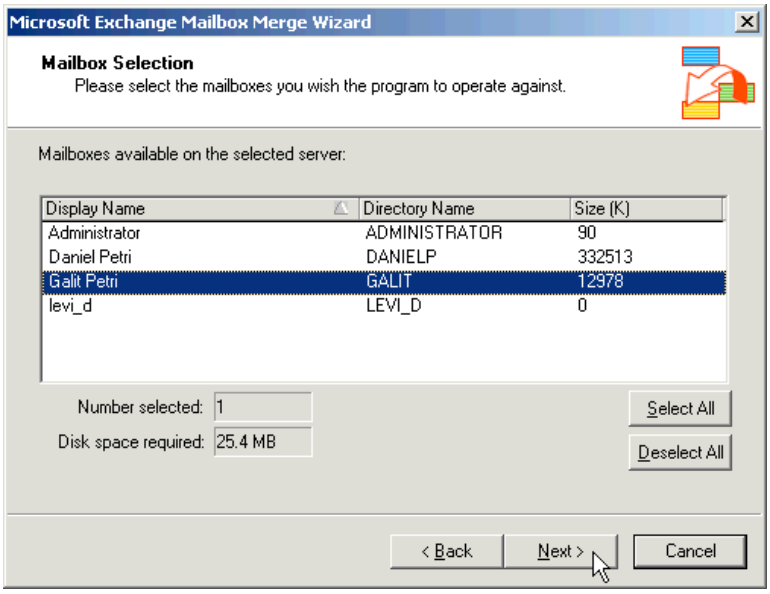

   Click Next.

8. On the "Locale Selection" dialogue box select the Locale that you would like to use to search the mailboxes.

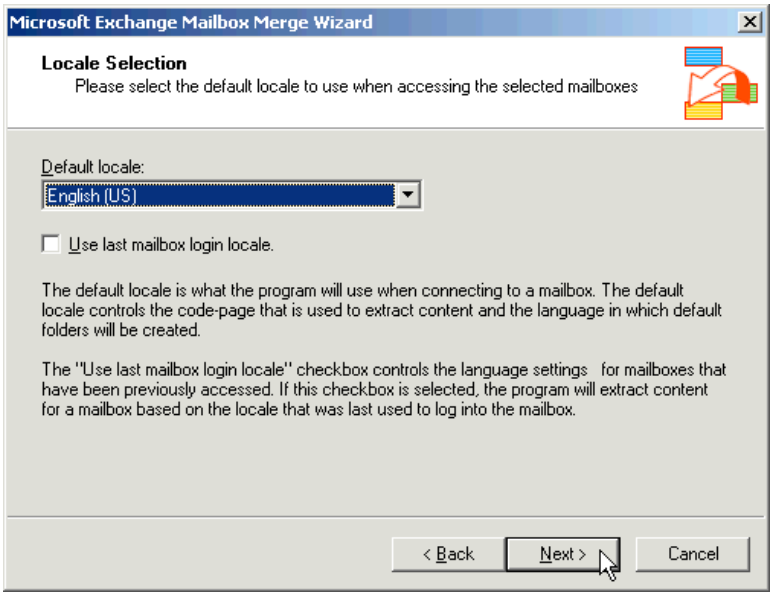
Microsoft Exchange Mailbox Merge Wizard
Locale Selection
Please select the default locale to use when accessing the selected mailboxes
Default locale:
English (US)
Use last mailbox login locale.
The default locale is what the program will use when connecting to a mailbox. The default locale controls the code-page that is used to extract content and the language in which default folders will be created.
The "Use last mailbox login locale" checkbox controls the language settings for mailboxes that have been previously accessed. If this checkbox is selected, the program will extract content for a mailbox based on the locale that was last used to log into the mailbox.
< Back
Next >
Cancel

Click Next.

9. Specify the path to the folder where you want to place the .PST files. This folder should reside on a different HD that the one where your mailbox stores are located, but this is only a recommendation for performance benefits, not a must. Also, make sure you have enough free space on that partition.

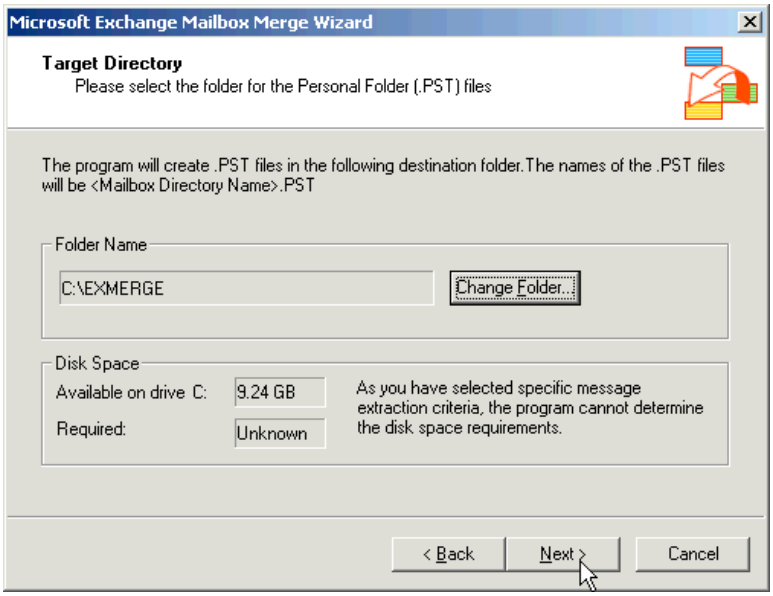
Microsoft Exchange Mailbox Merge Wizard
Target Directory
Please select the folder for the Personal Folder (.PST) files
The program will create .PST files in the following destination folder.The names of the .PST files will be <Mailbox Directory Name>.PST
Folder Name
C:\EXMERGE
Change Folder...
Disk Space
Available on drive C:
9.24 GB
Required:
Unknown
As you have selected specific message extraction criteria, the program cannot determine the disk space requirements.
< Back
Next >
Cancel

Click Next.

10. On the "Save Settings" dialogue box you can now save the settings you've configured so far, **DO NOT** click Next.

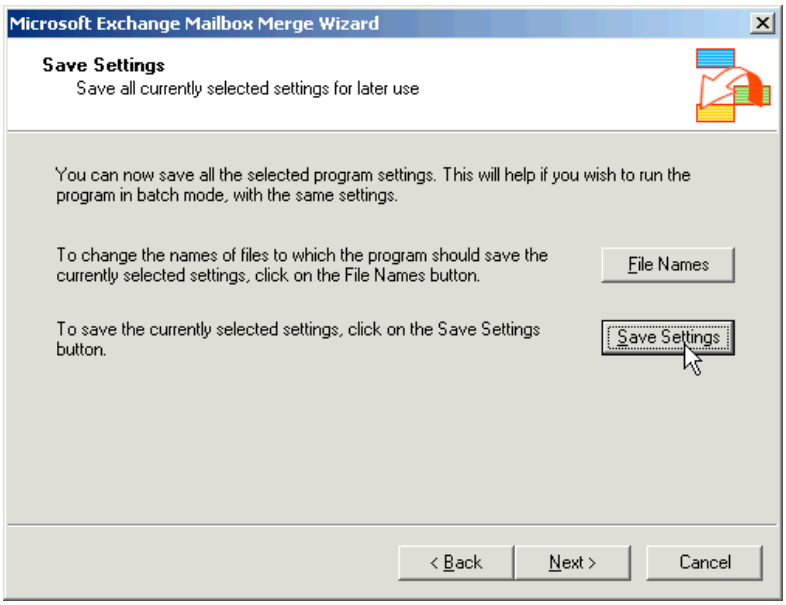
Microsoft Exchange Mailbox Merge Wizard
Save Settings
Save all currently selected settings for later use
You can now save all the selected program settings. This will help if you wish to run the program in batch mode, with the same settings.
To change the names of files to which the program should save the currently selected settings, click on the File Names button.
File Names
To save the currently selected settings, click on the Save Settings button.
Save Settings
< Back
Next >
Cancel

Make sure you save the EXMERGE configuration files in a convenient path, i.e. c:\EXMERGE\ or similar.

11. Quit EXMERGE
12. Notice that the saved files location (i.e. c:\EXMERGE\) holds a few files. At the minimum, it will hold the following files:

    EXMERGE.INI

    MAILBOXES.TXT
13. You can now run the EXMERGE operation via the command line or through a simple batch file

    EXMERGE -F C:\EXMERGE\EXMERGE.INI -B

    (use your own path to the EXMERGE.INI file)
14. Note how the .PST file(s) will be added to the folder that was specified as the file location. Note that this process might take a considerable amount of time, all depending on the number and size of the mailboxes being exported. You ill not get a visual confirmation of the operation, and you will not be notified when it has ended, however, if you're in doubt, use Task Manager to see if the EXMERGE.EXE process is still running.

If you want to get a visual GUI you can add -D to the command, thus causing EXMERGE to display the GUI while running:

EXMERGE -F C:\EXMERGE\EXMERGE.INI -B -D

Restoring the mailboxes

You have several options when you need to restore mail from a mailbox's .PST file. You can run EXMERGE in normal mode and select the .PST data by a range of dates or other criteria and restore them directly to any production mailbox.

Or you can copy the .PST file to the user's local disk or make it available on a network share. Then, instruct the user to add Personal Folders to their Outlook Profile and add the .PST file when they do so. The user can then browse through the .PST file's contents and retrieve the messages they are looking for.

Question: Can I export mailboxes straight out of Exchange and into the Snapshot EFB appliance?

Answer: Yes. The Snapshot EFB appliance has a Samba SMB service built into it that enables it to present Windows file shares onto the network. The Samba configuration file on the Snapshot EFB appliance is held in the directory /etc/samba/. By default Snapshot EFB have an ***exports*** share setup against the directory /mnt/i/a/snapshot.efb/exports for this very purpose. You can move this share to another drive and archive directory by modifying the Samba configuration file.

Question: If I export PST files directly to a share on the Snapshot EFB appliance can the PST files be synchronised to a caddy?

Answer: Yes. As long as you export the PST files to an export directory of a registered Archive then you can synchronise the PST files to a caddy just be assigning the archive to a caddy.

SQL Server Backup

The backup of an SQL server can be split into two distinct operations. The backup of the OS and SQL server database container file and the backup of the SQL database transactions.

Using the Snapshot file anti-locking agent Snapshot can backup the operating system and the files holding the SQL databases. This type of backup can be used to restore a server's operating system, applications and database structure. This is a full server file backup and is described above.

It may in some cases in maybe OK to backup an SQL database on an operational basis using a standard file backup, as long as the database is not being used at the time. However any changes to transactions within the database whilst the backup is occurring will cause the database file to become inconsistent and therefore unusable. To circumvent this Snapshot can extract transactions out of the SQL database and backup those transactions. This is called an SQL backup.

To get Snapshot to automate and perform SQL transaction extracts from SQL Server databases and back those SQL extracts follow the procedure below:

1. Create a machine group to hold the IP addresses of the SQL servers
2. Enter into the group the IP addresses of the SQL servers
3. Create a schedule to perform a backup of the SQL server group within the archive that is to hold the SQL data
    a. Make sure that the schedule uses the SQL group yu have just created
    b. Add in the following advanced options commands

    -8 'sql' -u 'mssqldmp' -b 'c:\tmp\SQL_DUMP' -b 'c:\etc\snapshot-efb.conf\log.d' -b 'c:\etc\snapshot-efb.conf\scripts'

-8 'sql'	This is a label
-u 'mssqldmp'	This is pre execution script that snapshot will run before doing a backup. This script runs an SQL script that creates the database export files.

-b 'c:\tmp\SQL_DUMP' This is the place where the SQL export files will be put. Actually at this point we are guessing where snapshot will put the files. Snapshot will choose the drive with the most free space to place the export file on.

c. Set the times you wish the schedule to run and save the schedule

4. Run the schedule
5. Go to the console menu
6. Select option 3 'UPDATE Snapshot EFB'
7. Select option 4 'AGENT Reg Export + MS SQL Check for Windows.'
8. Note the out put it will tell you what databases it has found and where it will but the export output for each server
9. Modify the SQL schedule you created to include the paths for each of the server SQL export positions. The paths will be consistent just the drive letters different, maybe. It does not matter that these paths will not exist on all the servers in the SQL group, Snapshot will just ignore if the path does not exist

By default Snapshot will setup an SQL backup schedule as follows:

Day 1 - FULL
Day 2 - Differential
Day 3 - Differential
Day 4 - Differential
Day 5 - Differential

Any additional backups done during the day will be performed as incremental.

Snapshot will perform it's own house keeping removing unneeded export files from the servers.

NOTE: FOR SQL SERVER 2005

Snapshot uses the command xp_cmdshell. To use xp_cmdshell whether you are a sysadmin or a regular user you need to first enable the use of xp_cmdshell. This can be done by using sp_configure:

```
sp_configure 'xp_cmdshell', '1'
```

Or by using the SQL Server Surface Area Configuration as follows:

Launch the SQL Server Surface Area Configuration.tool which can be found under SQL Server 2005 Configuration Tools in the SQL Server 2005 menus

Select "Surface Configuration Area for Features".

Check the box for "xp_cmdshell"

Select "Apply"

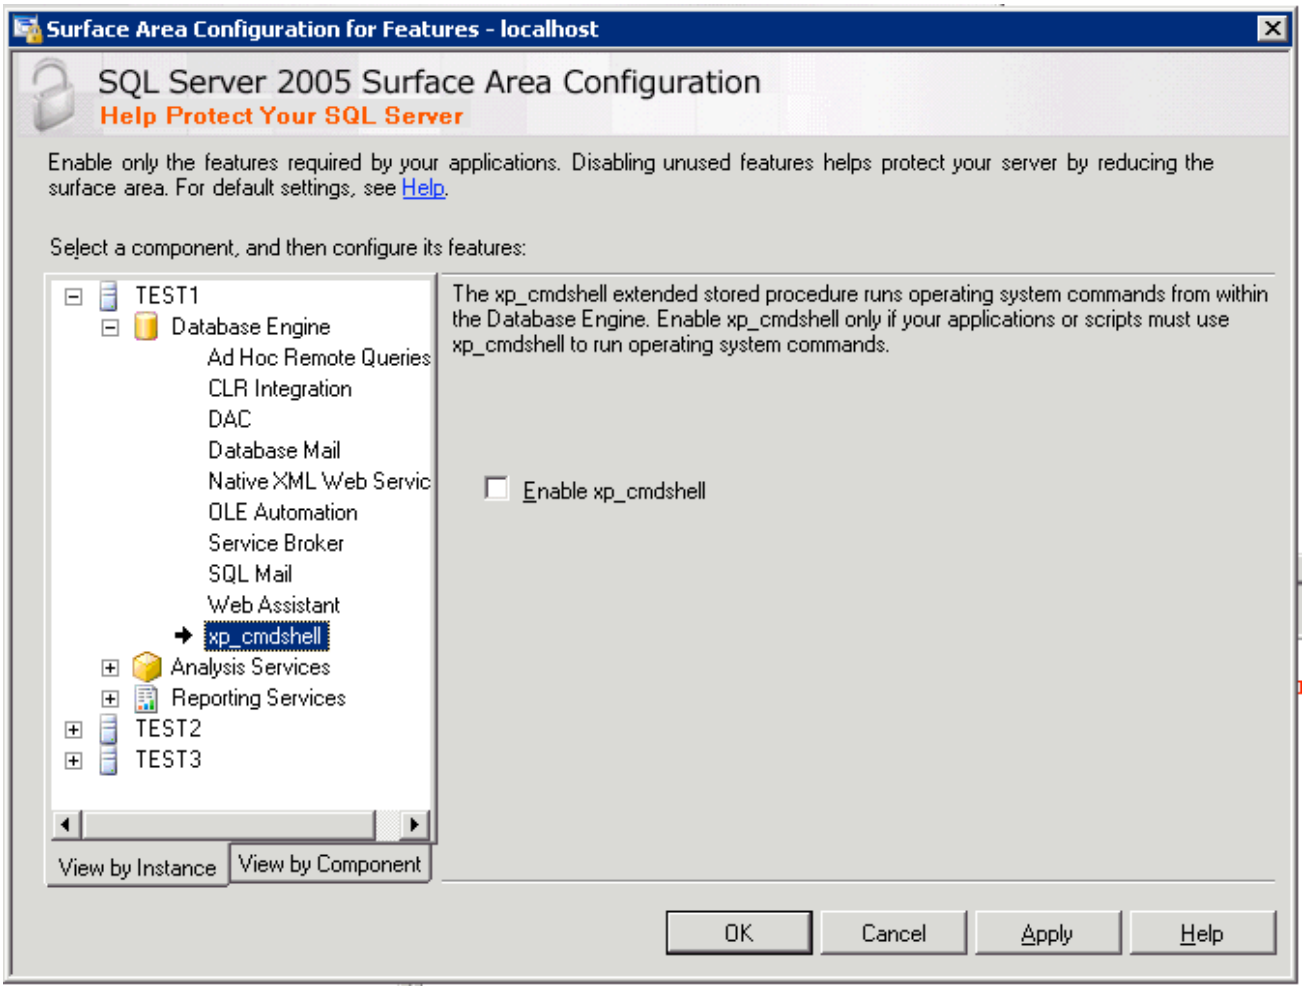
Surface Area Configuration for Features - localhost
SQL Server 2005 Surface Area Configuration
Help Protect Your SQL Server
Enable only the features required by your applications. Disabling unused features helps protect your server by reducing the surface area. For default settings, see Help.
Select a component, and then configure its features:
TEST1
Database Engine
Ad Hoc Remote Queries
CLR Integration
DAC
Database Mail
Native XML Web Servic
OLE Automation
Service Broker
SQL Mail
Web Assistant
xp_cmdshell
Analysis Services
Reporting Services
TEST2
TEST3
The xp_cmdshell extended stored procedure runs operating system commands from within the Database Engine. Enable xp_cmdshell only if your applications or scripts must use xp_cmdshell to run operating system commands.
Enable xp_cmdshell
View by Instance
View by Component
OK
Cancel
Apply
Help

Backup AD Database Backup

Snapshot anti-file locking will attempt to backup the AD database whilst it is mounted and running. However, this is not always successful. You can check whether Snapshot has automatically managed to backup the AD database by looking for a red skip on the AD machine backup and a bold entry for NTDS.DIT in the skipped files tab. If this occurs follow the instructions below.

Backup the AD database file NTDS.DIT using Windows backup. The windows backup natively supports backing up Active Directory whilst it is online. To backup the AD database run the Windows backup utility and select the option to back up System State in the backup wizard.

Backup the database into a directory, the default is the %SystemRoot%\ntds folder, ready for Snapshot to pickup as part of the normal daily file backup schedule.

AD Database Restore

If ntds.dit gets corrupted or deleted or is missing, this can happen if the promotion process to domain controller goes bad, you will have to manually recover it.

10. Recover a copy the ntds.dit file into the %SystemRoot%\ntds folder
11. Reboot the domain controller and press F8 to display the Windows 2000 Advanced Options menu.
12. Select Directory Services Restore Mode and then press ENTER.
13. Select the correct installation, and then press ENTER to start the boot process.
14. Logon using the administrator account and password you specified (if you ran Dcpromo.exe to install Active Directory, it requested a password to be used for the Administrator password for Active Directory Restore Mode. This password is not stored in Active Directory. It is stored in an NT4-style SAM file and is the only account available when the AD is corrupted)
15. Click OK. This acknowledges the warning message that you are using Safe mode
16. Click Start, Programs, Accessories, System Tools, and then click Backup
17. Select the Restore tab
18. Click the + symbol next to the following items to expand them:
    a. File
    b. Media Created
    c. System Drive
    d. Winnt
    e. NTDS
19. Click the NTDS folder to display the files in the folder
20. Click to select the ntds.dit check box
21. Leave the Restore files to box set to Original Location. This check box provides the option to restore to an alternative location
22. Click Start Restore

PUSH Synchronisation of Appliance to Appliance

Snapshot appliances can synchronise archives to another Snapshot appliance. This feature is commonly used to copy onsite data held in a local Snapshot Appliance to a central Snapshot Appliance. When setting up an appliance synchronisation attention must be paid to the size of the comms line and the amount of data flow between the appliances. Snapshot provides rich data flow information through the browser interface within the history tab. For new sites Snapshot can be used in Census mode to estimate the amount of daily data flow.

The following is a set of step by step instructions to setup a synchronisation between two Snapshot appliances where the Snapshot appliance backing up servers **pushes** data to the receiving appliance:

For the following instructions the:

Source appliance = The appliance backing up the servers
Target appliance = The appliance receiving the synchronisation data

1. From the Source Appliance - Setup automatic remote login to the target appliance
    a. Go to the console menu
    b. Select option 3
    c. Selection option 8
    d. Enter the IP address of the Target Appliance
    e. Enter the password for the Target Appliance (default password is snapshot but this can be changed via the appliance graphical start menu)
2. From the Source Appliance – Create a SYNC_PUSH group and add in the Target Appliance IP address to the group.

    NOTE: The SYNC_PUSH group can only be created manually

    a. Create a SYNC_PUSH group
        i. On the Source Appliance open a terminal window
        ii. cd /etc/snapshot-efb.conf/groups
        iii. touch SYNC_PUSH

iv. Login to the main server backup archive
v. Select the SETUP tab
vi. Select the SLAVE sub tab
vii. Select the SYNC_PUSH group and press the LOAD button
viii. Add in the IP address of the Target Appliance and save the group pressing the SAVE button

b. Now create a schedule using the SYNC_PUSH group to automate the synchronisation of the appliance
   i. Select the SETUP tab
   ii. Select the SCHEDULE tab
   iii. Select the SYNC_PUSH group
   iv. Select the time, days and months that you wish the synchronisation to happen
   v. Save the schedule

Snapshot will synchronise the source appliance to the target appliance according to the schedule on the source appliance. Data is pushed from the source appliance to the target appliance.

Q & A

Question: Can I exclude certain archives from being synchronised to the target appliance?

Answer: Yes. In the Advance Options section of the SYNC_PUSH schedule you can add excludes to exclude the synchronisation of archives.

For example to exclude an archive on drive 'd' in path /mnt/i/a/snapshot.efb the Advanced Option would be:

--exclude=d/snapshot.efb

You can have multiple excludes by adding many --exclude commands separated by spaces.

Question: How do I know when the SYNC_PUSH job is completed?

Answer: Snapshot EFB will send out an email when the SYNC_PUSH job is completed or errors informing you of the details of the SYNC_PUSH.

Question: I keep on getting emails saying that the SYNC_PUSH job has failed to start because the previous one is still running. Is this normal?

Answer: Depending on the amount of data being sent to the target Snapshot EFB appliance and the bandwidth of the communications line between the Source and Target Snapshot EFB appliances it may be that the previous SYNC_PUSH is still sending data. Snapshot EFB will not start a second SYNC_PUSH whilst an existing SYNC_PUSH job is still running. Simplistically the amount of time taken to complete a SYNC_PUSH is a product of the size of the network line and the amount of data be sent. Getting a SYNC_PUSH not started email is not a critical error, but if you are regularly getting them then you should investigate either cutting down the amount of data you are trying to sync or increasing the network bandwidth between the Source and Target Snapshot EFB appliances.

PULL Synchronisation of Appliance to Appliance

The basic intent is the same as the PUSH synchronisation, to copy archives from one Snapshot appliance to another. The difference from the PUSH synchronisation is that the PULL synchronisation is instigated from the target (receiving) Snapshot appliance doing a synchronisation to its caddies. Another difference is that a PULL synchronisation provides more control over which archives are synchronised. In a PULL synchronisation archives are individually selected for synchronisation and their destination position on the target appliance can be different from their position on the source appliance. Whereas a PUSH synchronisation synchronises the whole Snapshot appliance putting archives into the same position on the target Snapshot appliance.

1. Authenticate Target appliance to be able auto log into the Source appliance
    a. Using the Console Menu
    b. Select option 3
    c. Select option 8
    d. Press [return] to accept port 22
    e. Enter IP address of Source appliance
    f. Enter password for Source appliance
2. Setup receiving archive directory on Target appliance
    a. Open a terminal window – Start menu / Run
    b. Create the receiving archive directory - (e.g. If putting the archive on the 'b' drive:
    # mkdir /mnt/i/b/a
    # mkdir /mnt/i/b/a/snapshot.efb
    # mkdir /mnt/i/b/a/snapshot.efb/etc

    c. Copy the archive register.txt file from the Source appliance to the target machine. On the Target appliance:

    # scp [Source IP]:/mnt/i/[archive]/etc/register.txt /mnt/i/[archive]/etc

    (Eg. scp 192.168.1.6:/mnt/i/snapshot-efb/etc/register.txt /mnt/i/b/a/snapshot.efb/etc)

    d. Create the PULL sync file for the archive on the Target machine:

vi /mnt/i/b/a/snapshot.efb/etc/rsync_from.conf

Press the 'i' key to go into insert mode

Type in the follwing:

-e ssh [Source IP]:[archive directory]

(eg. –e ssh 192.168.1.4:/mnt/i/a/snapshot.efb)

Press the 'ESC' key

Press the keys ':wq'

3. On the Target appliance add the archive you have just created into the WEB browser interface

   Using a terminal open the Snapshot archive list file and add in the new target archive:

   # vi /etc/snapshot-efb.conf/snapshot.xxx

   Press the 'i' key to go into insert mode

   Add at bottom of list the new target archive. E.g.

   /mnt/i/b/a

   Press the 'ESC' key

   Press the keys ':wq'

   Refresh the browser interface on the target appliance. The new archive should now be in the archive drop down

4. Assign the archive to a caddy

   a. Option 4
   b. Option 6

Q & A

Question: When is a PULL rsync of an archive started?

Answer: A PULL sync of a remote archive is started after the archive on the target appliance has been synchronised to a caddy. This means for a PULL sync of an archive to happen the archive on the target appliance MUST be assigned to a caddy

Question: How do I tell if a PULL sync of an archive is running?

Answer: In a terminal on the target appliance run the following command. It will show you if the pull rsync process is still running

ps –ef | grep *[enter IP address of remote appliance]*

You can also run the following command a few times to see if disc space is being used. This will prove that data is being transferred

df

Question: Where can I see a log of the PULL sync to check for any errors?

Answer: On the target appliance look in *tmp* the directory under the archive you are synchronising to. In there will be a rsync.log file. Check the date of this file to see when it was last written to, this would show you when the PULL sync was either last run

eg. cd /mnt/i/a/snapshot.efb/tmp

ls –l

cat rsync.log

Question: What are the most common causes of a PULL remote sync not running?

Answer:

- The archive not assigned to a caddy on the target appliance
- The automated remote login for ssh not setup between the source and target appliances
- The target appliance not being able to see the source appliance on the network

Installation of Snapshot EFB eMail Archiving

On the Exchange server

Setup Microsoft Exchange Journaling

1. Login to the Exchange server as an 'administrator'
2. Create a user called 'mail-efb' with first name, login name and mail alias all with same names
3. Add a password for 'mail-efb' and note it down for later Snapshot EFB configuration
4. Execute Start->All Programs->Microsoft Exchange->System Manager
5. Right click on the mail-box store and enable message archiving as below:

   Servers->Storage Group->Mailbox (Right Click for properties)

   In the General Tab tick "Archive All Messages" and then browse for "mail-efb" mailbox.

   NOTE: You need to do this for each mail store object.

6. Make sure the IMAP service is running on the Exchange server

   Start IMAP under Control Panel Services
   Start IMAP in the Exchange Server Manager

Prerequisites:

Exchange 2003 with SP1 preferably latest service pack Journaling is supported from 2000 onwards.

Performance Impact:

Exchange will copy every email sent and received to the journal mailbox and will slow the server by about 30% and use double the disk space until the journal account is cleared every minute or so. Under normal circumstances very little extra space is used as the journal account is cleared constantly by Snapshot when setup.

On the Snapshot EFB appliance

1. Create a "m" archive on a free drive eg: /mnt/i/d

   mkdir /mnt/i/d/m

2. Snapshot needs to always think that the "m" archive is on the "a" drive so if you have created it on any other drive than the "a" drive then execute the following command

   ln -s /mnt/i/d/m /mnt/i/a/m

3. Download mailefb.tgz and thinderbird.tgz and extract into /mnt/i/a/m

   cd /mnt/i/a/m
   tar xvzf /tmp/thunderbird.tgz
   tar xvzf /tmp/mailefb.tgz
   cp -a TEMPLATE/snapshot.efb ./

4. Move to snapshot/etc and run setup scrips to set passwords

   cd snapshot.efb/etc
   vi fetchmail.conf

   Now you are editing the Snapshot-EFB mail configuration file that controls the collection of mail from the Exchange server. Set the following parameters:

   IP adress of Exchange server
   Login name for Exchange eg: mail-efb
   Password for the login name above on Exchange

5. Now run the setup script

   ./mkcert.sh

6. Now setup login password for IMAP via ssh on port 993

   ./setup_passwd.sh

7. Automate the email download from Exchange to the Snapshot-EFB appliance by adding an entry in the appliances scheduler

   Run the command - vi /etc/crontab

   Add the following line:

   * * * * * root /mnt/i/a/m/snapshot.efb/etc/start_fetchmail.sh

   You can manually start the email downloading from Exchange by running the command:

   /mnt/i/a/m/snapshot.efb/etc/start_fetchmail.sh

8. Once you have started the mail sync it is a good idea to check that it is running OK. To do this check the log files in the following directory

9. Add thunderbird to Graphical desktop menus. Add a menu option under "Programs"

   vi /root/.icewm/menu
   prog "Thunderbird Mail Client" - /bin/sh -c "/mnt/i/a/m/thunderbird.sh"

10. Run thunderbird from menu

    Add IMAP account
    Email address snapshot-efb@localdomain.com
    Incomming Server: 127.0.0.1
    Outgoing Server: 127.0.0.1
    Enable SSL encryption i.e. port 993 on incoming
    Login as "snapshot-efb" with the password you setup locally

Now add the eMail archive to the Snapshot EFB browser interface so that it can be licensed and assigned to a caddy

a. Add the eMail archive to the snapshot.xxx file

cd /etc/snapshot-efb.conf
vi snapshot.xxx

Press the 'i' key to go into insert mode
Add the line ...

/mnt/i/a/m

Press the 'ESC' key to exit insert mode
Press the 'wq' keys and 'Return' to save the file and exit the vi editor

b. Using the Firefox browser interface log into the eMail archive and add in the eMail archive license key in the Setup->Archive tab
c. Using the white console menu assign the eMail archive to a caddy for synchronisation